arXiv:0905.2517v2 [hep-ph] 18 Nov 2009

The GENIE Neutrino Monte Carlo Generator

C.Andreopoulos[a], A.Bell[b], D.Bhattacharya[b], F.Cavanna[c], J.Dobson[d], S.Dytman[b], H.Gallagher[e], P.Guzowski[d], R.Hatcher[f], P.Kehayias[e], A.Meregaglia[g,h], D.Naples[b], G.Pearce[a], A.Rubbia[g], M.Whalley[i], T.Yang[j]

[a] *Rutherford Appleton Laboratory, STFC, Oxfordshire OX11 0QX, UK*

[b] *Physics Dept., Pittsburgh University, Pittsburgh PA 15260, USA*

[c] *Physics Dept., L'Aquila University, 67100 L'Aquila, Italy*

[d] *Physics Dept., Imperial College London, Blackett Lab., London SW7 2BW, UK*

[e] *Physics Dept., Tufts University, Medford, MA, 02155, USA*

[f] *Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

[g] *Physics Dept., ETH Zurich, CH-8093 Zurich, Switzerland*

[h] *IPHC - Strasbourg, F-67037 Strasbourg Cedex 2, France*

[i] *Physics Dept., Durham University, Durham DH1 3LE, UK*

[j] *Physics Dept., Stanford University, Stanford, CA 94309, USA*

Abstract

GENIE [1] is a new neutrino event generator for the experimental neutrino physics community. The goal of the project is to develop a 'canonical' neutrino interaction physics Monte Carlo whose validity extends to all nuclear targets and neutrino flavors from MeV to PeV energy scales. Currently, emphasis is on the few-GeV energy range, the challenging boundary between the non-perturbative and perturbative regimes, which is relevant for the current and near future long-baseline precision neutrino experiments using accelerator-made beams. The design of the package addresses many challenges unique to neutrino simulations and supports the full life-cycle of simulation and generator-related analysis tasks.

GENIE is a large-scale software system, consisting of ∼120,000 lines of C++ code, featuring a modern object-oriented design and extensively validated physics content. The first official physics release of GENIE was made available in August 2007, and at the time of the writing of this article, the latest available version was v2.4.4.

Key words: Neutrino; Monte Carlo Generator; Neutrino Interaction; Neutrino-induced Hadronization; Intra-nuclear Hadron Transport; Rescattering; GENIE; AGKY; INTRANUKE

1. Introduction

Over the last few years, throughout the field of high energy physics (HEP), we have witnessed an enormous effort committed to migrating many popular procedural Monte Carlo Generators into their C++ equivalents designed using object-oriented methodologies. Well-known examples are the GEANT [2], HERWIG [3] and PYTHIA [4] Monte Carlo Generators. This reflects a radical change in our approach to scientific computing. Along with the eternal requirement that the modeled physics be correct and extensively validated with external data, the evolving nature of computing in HEP has introduced new requirements. These requirements relate to the way large HEP software systems are developed and maintained, by wide geographically-spread

collaborations over a typical time-span of ∼25 years during which they will undergo many (initially unforeseen) extensions and modifications to accommodate new needs. This puts a stress on software qualities such as re-usability, maintainability, robustness, modularity and extensibility. Software engineering provides many well proven techniques to address these requirements and thereby improves the quality and lifetime of HEP software. In neutrino physics, the requirements for a new physics generator are more challenging for three reasons: the lack of a 'canonical' procedural generator, theoretical and phenomenological challenges in modeling few-GeV neutrino interactions, and the rapidly evolving experimental and theoretical landscape.

Neutrinos come from many sources and a variety of experiments have been mounted to measure their properties. These experiments have complicated detectors composed of many elements and the neutrinos have many flavors and a wide energy spectrum (from $\sim$ 1 MeV to $\sim$ 1 PeV). Our long-term goal is for GENIE to become a 'canonical' neutrino event generator with wide applicability. The origins of the code come from the Soudan experiment [5] and recent application has been primarily to MINOS [6]. Thus, emphasis has been given to the few-GeV energy range, the challenging boundary between the non-perturbative and perturbative regimes. These are relevant for current and near future long-baseline precision neutrino oscillation experiments using accelerator-made beams, one of the focuses of high energy physics. GENIE development over the next five years will be driven by the upcoming generation of accelerator experiments including T2K [7], NoVA [8], Minerva [9] MicroBooNE [10] and ArgoNEUT [11]. These developments are well underway and the code is being used successfully in each of these experiments. The present version provides comprehensive neutrino interaction modelling in the energy from ∼100 MeV to a few hundred GeV. Results can be obtained and will be qualitatively correct for any nuclear target.

GENIE[1] is a ROOT-based [12] Neutrino MC Generator. It was designed using object-oriented methodologies and developed entirely in C++ over a period of more than three years, from 2004 to 2007. Its first official physics release (v2.0.0) was made available on August 2007 and, at the time of writing this article, the latest available version was v2.4.4. It also describes v2.6.0 which will be released shortly. GENIE has already been adopted by the majority of neutrino experiments, including those using the JPARC and NuMI neutrino beamlines, and will be an important physics tool for the exploitation of the world accelerator neutrino program.

The project is supported by a group of physicists from all major neutrino experiments operating in this energy range, establishing GENIE as a major HEP event generator collaboration. Many members of the GENIE collaboration have extensive experience in developing and maintaining the legacy Monte Carlo Generators that GENIE strives to replace, which guarantees knowledge exchange and continuation. The default set of physics models in GENIE have adiabatically evolved from those in the NEUGEN [13] package, which was used as the event generator by numerous experiments over the past decade.

This article will discuss the paradigm shift brought about by GENIE in neutrino physics simulations. In Sec. 2 we describe the unique challenges facing neutrino simulations in more detail. Section 3 gives a brief overview of the physics models available in GENIE. Section 4 gives a brief discussion of upgrades in progress. Section 5 describes the object-oriented design of GENIE. Section 6 describes the GENIE applications and utilities available for simulation and analysis tasks. Section 7 describes the structure of the GENIE collaboration and Section 8 describes code availability, distribution, supported platforms, external dependencies, releases, and license.

[1] *GENIE* stands for *G*enerates *E*vents for *N*eutrino *I*nteraction *E*xperiments

2. Neutrino Interaction Simulation: Challenges and Significance

Neutrinos have played an important role in particle physics since their discovery half a century ago. They have been used to elucidate the structure of the electroweak symmetry groups, to illuminate the quark nature of hadrons, and to confirm our models of astrophysical phenomena. With the discovery of neutrino oscillations using atmospheric, solar, accelerator, and reactor neutrinos, these elusive particles now take center stage as the objects of study themselves. Precision measurements of the lepton mixing matrix, the search for lepton charge-parity (CP) violation, and the determination of the neutrino masses and hierarchy will be major efforts in HEP for several decades. The cost of this next generation of experiments will be significant, typically tens to hundreds of millions of dollars. A comprehensive, thoroughly tested neutrino event generator package plays an important role in the design and execution of these experiments, since this tool is used to evaluate the feasibility of proposed projects and estimate their physics impact, make decisions about detector design and optimization, analyze the collected data samples, and evaluate systematic errors. With the advent of high-intensity neutrino beams from proton colliders, we have entered a new era of high-statistics, precision neutrino experiments which will require a new level of accuracy in our knowledge, and simulation, of neutrino interaction physics [14].

While object-oriented physics generators in other fields of high energy physics were evolved from well established legacy systems, in neutrino physics no such 'canonical' MC exists. Until quite recently, most neutrino experiments developed their own neutrino event generators. This was due partly to the wide variety of energies, nuclear targets, detectors, and physics topics being simulated. Without doubt these generators, the most commonly used of which have been GENEVE [15], NEUT [16], NeuGEN [13], NUANCE [17] and NUX [18], played an important role in the design and exploration of the previous and current generation of accelerator neutrino experiments. Tuning on the basis of unpublished data from each group's own experiment has not been unusual making it virtually impossible to perform a global, independent evaluation for the state-of-the-art in neutrino interaction physics simulations. Moreover, limited manpower and the fragility of the overextended software architectures meant that many of these legacy physics generators were not keeping up with the latest theoretical ideas and experimental measurements. A more recent development in the field has been the direct involvement of theory groups in the development of neutrino event generators, such as the NuWRO [19] and GiBUU [20] packages, and the inclusion of neutrino scattering in the long-established FLUKA hadron scattering package [21].

Simulating neutrino interactions in the energy range of interest to current and near-future experiments poses significant challenges. This broad energy range bridges the perturbative and non-perturbative pictures of the nucleon and a variety of scattering mechanisms are important. In many areas, including elementary cross sections, hadronization models, and nuclear physics, one is required to piece together models with different ranges of validity in order to generate events over all of the available phase space. This inevitably introduces challenges in merging and tuning models, making sure that double counting and discontinuities are avoided. In addition there are kinematic regimes which are outside the stated range of validity of all available models, in which case we are left with the challenge of developing our own models or deciding which model best extrapolates into this region. An additional fundamental problem in this energy range is a lack of data. Most simulations have been tuned to bubble chamber data taken in the 70's and 80's. Because of the limited size of the data samples (important exclusive channels might only contain a handful of events), and the limited coverage in the space of ($\nu/\overline{\nu}, E_\nu$, A), substantial uncertainties exist in numerous aspects of the simulations.

The use cases for GENIE are also informed by the experiences of the developers and users of the previous generation of procedural codes. Dealing with these substantial model uncertainties has been an important analysis challenge for many recent experiments. The impact of these uncertainties on physics analyses have been evaluated in detailed systematics studies and in some cases the models have been fit directly to experimental data to reduce systematics. These

'downstream' simulation-related studies can often be among the most challenging and time-consuming in an analysis.

To see the difficulties facing the current generation of neutrino experiments, one can look no further than the K2K and MiniBooNE experiments. Both of these experiments have measured a substantially different Q^2 distribution for their quasielastic-like events when compared with their simulations, which involve a standard Fermi Gas model nuclear model [22, 23]. The disagreement between nominal Monte Carlo and data is quite large - in the lowest Q^2 bin of MiniBooNE the deficit in the data is around 30% [23]. It seems likely that the discrepancies seen by both experiments have a common origin. However the two experiments have been able to obtain internal consistency with very different model changes - the K2K experiment does this by eliminating the charged current (CC) coherent contribution in the Monte Carlo [24] and the MiniBooNE experiment does this by modifying certain parameters in their Fermi Gas model [23]. Another example of the rapidly evolving nature of this field is the recently reported excess of low energy electron-like events by the MiniBooNE collaboration [25]. These discrepancies have generated significant new theoretical work on these topics over the past several years [26, 27, 28, 29, 30, 31, 32, 33]. The situation is bound to become even more interesting, and complicated, in the coming decade, as new high-statistics experiments begin taking data in this energy range. Designing a software system that can be responsive to this rapidly evolving experimental and theoretical landscape is a major challenge.

In this paper we will describe the ways in which the GENIE neutrino event generator addresses these challenges. These solutions rely heavily on the power of modern software engineering, particularly the extensibility, modularity, and flexibility of object oriented design, as well as the combined expertise and experience of the collaboration with previous procedural codes.

3. Neutrino Interaction Physics models in GENIE

The set of physics models used in GENIE incorporates the dominant scattering mechanisms from several MeV to several hundred GeV and are appropriate for any neutrino flavor and target type. Over this energy range, many different physical processes are important. These physics models can be broadly categorized into nuclear physics models, cross section models, and hadronization models.

The neutrino-nucleus interaction involves a large variety of processes, all of which must be modeled to get an accurate description of the experimental signature of any detector and its many components. Since most theoretical models describe a small subset of these processes, GENIE must include many models. The broad energy range and the many nuclei to be covered force choosing models that have very broad applicability.

The particle in the nucleus with which the neutrino interacts depends strongly on the energy. At high energies (E_ν >10 GeV) neutrinos interact with a single quark inside a nucleon (neutron or proton); the code must model this interaction and the distribution of the residual quarks. At lower energies, the struck particles are neutrons and protons. The neutrino tends to strike a single nucleon (impulse approximation) which is affected by the nuclear medium in which it resides. In the high energy regime, the large body of neutrino-nucleon data is sufficient for development of a full model. At lower energies, neutrino-nucleon data has been used for the basic process and nuclear models developed for other probes (especially electrons) are adopted.

The 2 recent major developments have been in the transition region and the final state interactions (FSI) model. In physics model development for GENIE we have been forced to pay particular attention to this 'transition region', as for few-GeV experiments it dominates the event rate. In particular the boundaries between regions of validity of different models need to be treated with care in order to avoid theoretical inconsistencies, discontinuities in generated distributions, and double-counting. Treatment of FSI involves many aspects of nuclear physics and strong

interactions. Many events where particles are produced by the neutrinos have their topologies and kinematics altered. There are many effects to include and some dispute about the right techniques to use. Therefore, FSI treatment is one of the largest differences among models of the neutrino-nucleus interaction.

In this brief section we will describe the models available in GENIE and the ways in which we combine models to cover regions of phase space where clear theoretical or empirical guidance is lacking.

3.1. Nuclear Physics Model

The importance of the nuclear model depends strongly on the kinematics of the reaction. Nuclear physics plays a large role in every aspect of neutrino scattering simulations at few-GeV energies and introduces coupling between several aspects of the simulation. The relativistic Fermi gas (RFG) nuclear model is used for all processes. GENIE uses the version of Bodek and Ritchie which has been modified to incorporate short range nucleon-nucleon correlations [34]. This is simple, yet applicable across a broad range of target atoms and neutrino energies. The best tests of the RFG model come from electron scattering experiments [35]. At high energies, the nuclear model requires broad features due to shadowing and similar effects. At the lower end of the GENIE energy range, the impulse approximation works very well and the RFG is often successful. The nuclear medium gives the struck nucleon a momentum and average binding energy which have been determined in electron scattering experiments. Mass densities are taken from review articles [36]. For A <20, the modified Gaussian density parameterization is used. For heavier nuclei, the 2-parameter Woods-Saxon density function is used. Thus, the model can be used for *all* nuclei. Presently, fit parameters for selected nuclei are used with interpolations for other nuclei. All isotopes of a particular nucleus are assumed to have the same density.

It is well known that scattering kinematics for nucleons in a nuclear environment are different from those obtained in scattering from free nucleons. For quasi-elastic and elastic scattering, Pauli blocking is applied as described in Sec. 3.2. For nuclear targets a nuclear modification factor is included to account for observed differences between nuclear and free nucleon structure functions which include shadowing, anti-shadowing, and the EMC effect [37].

Nuclear reinteractions of produced hadrons is simulated using a cascade Monte Carlo which will be described in more detail in a following section. The struck nucleus is undoubtedly left in a highly excited state and will typically de-excite by emitting nuclear fission fragments, nucleons, and photons. At present de-excitation photon emission is simulated only for oxygen [38, 39] due to the significance of these 3-10 MeV photons in energy reconstruction at water Cherenkov detectors. Future versions of the generator will handle de-excitation photon emission from additional nuclear targets.

3.2. Cross section model

The cross section model provides the calculation of the differential and total cross sections. During event generation the total cross section is used together with the flux to determine the energies of interacting neutrinos. The cross sections for specific processes are then used to determine which interaction type occurs, and the differential distributions for that interaction model are used to determine the event kinematics. While the differential distributions must be calculated event-by-event, the total cross sections can be pre-calculated and stored for use by many jobs sharing the same physics models. Over this energy range neutrinos can scatter off a variety of different 'targets' including the nucleus (via coherent scattering), individual nucleons, quarks within the nucleons, and atomic electrons.

Quasi-Elastic Scattering: Quasi-elastic scattering (e.g. $\nu_\mu + n \rightarrow \mu^- + p$) is modeled using an implementation of the Llewellyn-Smith model [40]. In this model the hadronic weak current is expressed in terms of the most general Lorentz-invariant form factors. Two are set to zero as they

violate G-parity. Two vector form factors can be related via CVC to electromagnetic form factors which are measured over a broad range of kinematics in electron elastic scattering experiments. Several different parametrizations of these electromagnetic form factors including Sachs [41], BBA2003 [42] and BBBA2005 [43] models are available with BBBA2005 being the default. Two form factors - the pseudo-scalar and axial vector, remain. The pseudo-scalar form factor is assumed to have the form suggested by the partially conserved axial current (PCAC) hypothesis [40], which leaves the axial form factor $F_A(Q^2)$ as the sole remaining unknown quantity. $F_A(0)$ is well known from measurements of neutron beta decay and the Q^2 dependence of this form factor can only be determined in neutrino experiments and has been the focus of a large amount of experimental work over several decades. In GENIE a dipole form is assumed, with the axial vector mass m_A remaining as the sole free parameter with a default value of 0.99 GeV/c^2.

For nuclear targets, the struck a suppression factor is included from an analytic calculation of the rejection factor in the Fermi Gas model, based on the simple requirement that the momentum of the outgoing nucleon exceed the fermi momentum k_F for the nucleus in question. Typical values of k_F are 0.221 GeV/c for nucleons in ^{12}C, 0.251 GeV/c for protons in ^{56}Fe, and 0.256 GeV/c for neutrons in ^{56}Fe.

Elastic Neutral Current Scattering: Elastic neutral current processes are computed according to the model described by Ahrens et al. [44], where the axial form factor is given by:

$$G_A(Q^2) = \frac{1}{2}\frac{G_A(0)}{(1+Q^2/M_A^2)^2}(1+\eta). \quad (1)$$

The adjustable parameter η includes possible isoscalar contributions to the axial current, and the GENIE default value is $\eta = 0.12$. For nuclear targets the same reduction factor described above is used.

Baryon Resonance Production: The production of baryon resonances in neutral and charged current channels is included with the Rein-Sehgal model [45]. This model employs the Feynman-Kislinger-Ravndal [46] model of baryon resonances, which gives wavefunctions for the resonances as excited states of a 3-quark system in a relativistic harmonic oscillator potential with spin-flavor symmetry. In the Rein-Sehgal paper the helicity amplitudes for the FKR model are computed and used to construct the cross sections for neutrino-production of the baryon resonances. From the 18 resonances of the original paper we include the 16 that are listed as unambiguous at the latest PDG baryon tables and all resonance parameters have been updated. In our implementation of the Rein-Sehgal model interference between neighboring resonances has been ignored. Lepton mass terms are not included in the calculation of the differential cross section, but the effect of the lepton mass on the phase space boundaries is taken into account. For tau neutrino charged current interactions an overall correction factor to the total cross section is applied to account for neglected elements (pseudoscalar form factors and lepton mass terms) in the original model. The default value for the resonance axial vector mass m_A is 1.12 GeV/c^2, as determined in the global fits carried out in Reference [47].

Coherent Neutrino-Nucleus Scattering: Coherent scattering results in the production of forward going pions in both charged current ($\nu_\mu + A \rightarrow \mu^- + \pi^+ + A$) and neutral current ($\nu_\mu + A \rightarrow \nu_\mu + \pi^0 + A$) channels. Coherent neutrino-nucleus interactions are modeled according to the Rein-Sehgal model [48]. Since the coherence condition requires a small momentum transfer to the target nucleus, it is a low-Q^2 process which is related via PCAC to the pion field. The Rein-Sehgal model begins from the PCAC form at Q^2=0, assumes a dipole dependence for non-zero Q^2, with $m_A = 1.00$ GeV/c^2, and then calculates the relevant pion-nucleus cross section from measured data on total and inelastic pion scattering from protons and deuterium [49]. The GENIE implementation is using the modified PCAC formula described in a recent revision of the Rein-Sehgal model [50] that includes lepton mass terms.

Non-Resonance Inelastic Scattering: Deep (and not-so-deep) inelastic scattering (DIS) is calculated in an effective leading order model using the modifications suggested by Bodek and

Yang [37] to describe scattering at low Q^2. In this model higher twist and target mass corrections are accounted for through the use of a new scaling variable and modifications to the low Q^2 parton distributions. The cross sections are computed at a fully partonic level (the $\nu q \rightarrow l q'$ cross sections are computed for all relevant sea and valence quarks). The longitudinal structure function is taken into account using the Whitlow R ($R = F_L/2xF_1$) parameterization [51]. The default parameter values are those given in [37], which are determined based on the GRV98 LO parton distributions [52]. An overall scale factor of 1.032 is applied to the predictions of the Bodek-Yang model to achieve agreement with the measured value of the neutrino cross section at high energy (100 GeV). This factor is necessary since the Bodek-Yang model treats axial and vector form modifications identically and would therefore not be expected to reproduce neutrino data perfectly. This overall DIS scale factor needs to be recalculated when elements of the cross section model are changed.

The same model can be extended to low energies; it is the model used for the nonresonant processes that compete with resonances in the few-GeV region.

Quasi-Elastic Charm Production: QEL charm production is modeled according to the Kovalenko local duality inspired model [53] tuned by the GENIE authors to recent NOMAD data [54].

Deep-Inelastic Charm Production: DIS charm production is modeled according to the Aivazis, Olness and Tung model [55]. Charm-production fractions for neutrino interactions are taken from [56], and utilize both Peterson [57] and Collins-Spiller [58] fragmentation functions, with Peterson fragmentation functions being the default. The charm mass is adjustable and is set to 1.43 GeV/c^2 by default.

Inclusive Inverse Muon Decay: Inclusive Inverse Muon Decay cross section is computed using an implementation of the Bardin and Dokuchaeva model [59] taking into account all 1-loop radiative corrections.

Neutrino-Electron Elastic Scattering: The cross sections for all $\nu e-$ scattering channels other than Inverse Muon Decay is computed according to [60]. Inverse Tau decay is neglected.

3.2.1. Modeling the transition region

As discussed, for example, by Kuzmin, Lyubushkin and Naumov [61] one typically considers the total νN CC scattering cross section as

$$\sigma^{tot} = \sigma^{QEL} \oplus \sigma^{1\pi} \oplus \sigma^{2\pi} \oplus ... \oplus \sigma^{1K} \oplus ... \oplus \sigma^{DIS}$$

In the absence of a model for exclusive inelastic multi-particle neutrinoproduction, the above is usually being approximated as

$$\sigma^{tot} = \sigma^{QEL} \oplus \sigma^{RES} \oplus \sigma^{DIS}$$

assuming that all exclusive low multiplicity inelastic reactions proceed primarily through resonance neutrinoproduction. For the sake of simplicity, small contributions to the total cross section in the few GeV energy range, such as coherent and elastic νe^- scattering, were omitted from the expression above. In this picture, one should be careful in avoiding double counting the low multiplicity inelastic reaction cross sections.

In GENIE release the total cross sections is constructed along the same lines, adopting the procedure developed in NeuGEN [13] to avoid double counting. The total inelastic differential cross section is computed as

$$\frac{d^2\sigma^{inel}}{dQ^2 dW} = \frac{d^2\sigma^{RES}}{dQ^2 dW} + \frac{d^2\sigma^{DIS}}{dQ^2 dW}$$

The term $d^2\sigma^{RES}/dQ^2dW$ represents the contribution from all low multiplicity inelastic channels proceeding via resonance production. This term is computed as

$$\frac{d^2\sigma^{RES}}{dQ^2dW} = \sum_k \left(\frac{d^2\sigma^{R/S}}{dQ^2dW}\right)_k \cdot \Theta(Wcut - W)$$

where the index k runs over all baryon resonances taken into account, W_{cut} is a configurable parameter and $(d^2\sigma^{RS}_{\nu N}/dQ^2dW)_k$ is the Rein-Seghal model prediction for the k^{th} resonance cross section.

The DIS term of the inelastic differential cross section is expressed in terms of the differential cross section predicted by the Bodek-Yang model appropriately modulated in the "resonance-dominance" region $W < W_{cut}$ so that the RES/DIS mixture in this region agrees with inclusive cross section data [62, 63, 64, 65, 66, 67, 68, 69, 70, 71] and exclusive 1-pion [72, 73, 74, 75, 76, 77, 78, 79, 80, 69, 81] and 2-pion [82, 76] cross section data:

$$\begin{aligned}\frac{d^2\sigma^{DIS}}{dQ^2dW} &= \frac{d^2\sigma^{DIS,BY}}{dQ^2dW} \cdot \Theta(W - Wcut) + \\ &+ \frac{d^2\sigma^{DIS,BY}}{dQ^2dW} \cdot \Theta(Wcut - W) \cdot \sum_m f_m\end{aligned}$$

In the above expression, m refers to the multiplicity of the hadronic system and, therefore, the factor f_m relates the total calculated DIS cross section to the DIS contribution to this particular multiplicity channel. These factors are computed as $f_m = R_m \cdot P^{had}_m$ where R_m is a tunable parameter and P^{had}_m is the probability, taken from the hadronization model, that the DIS final state hadronic system multiplicity would be equal to m. The approach described above couples the GENIE cross section and hadronic multiparticle production model [83].

3.3. Neutrino-induced hadronic multiparticle production modeling

Neutrino-induced hadronic shower modeling is an important aspect of the intermediate energy neutrino experiment simulations, as non-resonant inelastic scattering becomes the dominant interaction channel for neutrino energies as low as 1.5 GeV.

Experiments are sensitive to the details of hadronic system modeling in many different ways. Experiments making calorimetric measurements of neutrino energy in charged current reactions are typically calibrated using single particle test beams, which introduces a model dependence in determining the conversion between detector activity and the energy of neutrino-produced hadronic systems. Physics analysis can also depend on the prediction of the hadron shower characteristics, such as shower shapes, energy profile and particle content, primarily for event identification. A characteristic example is a $\nu_\mu \rightarrow \nu_e$ appearance analysis, where the evaluation of backgrounds coming from neutral current (NC) events, would be quite sensitive on the details of the NC shower simulation and specifically the π^0 shower content. It is therefore imperative that the state-of-the-art in shower modeling is included in our neutrino interaction simulations.

GENIE uses the AGKY hadronization model which was developed for the MINOS experiment [83]. This model integrates an empirical low-invariant mass model with PYTHIA-6 at higher invariant masses. The transition between these two models takes place over an adjustable window with the default values of 2.3 GeV/c^2 to 3.0 GeV/c^2, so as to ensure continuity of all simulated observables as a function of the invariant mass. For the hadronization of low-mass states the model proceeds in two phases, first determining the particle content of the hadronic shower, and secondly determining the 4-momenta of the produced particles in the hadronic center of mass.

The AGKY's low mass hadronization model generates hadronic systems that typically consist of exactly one baryon (p or n) and any number of π^+, π^-, π^0, K^+, K^-, K^0, $\bar{K^0}$ mesons kinematically possible and allowed by charge conservation.

For a fixed hadronic invariant mass and initial state (neutrino and hit nucleon), the algorithm for generating the hadron shower particles generally proceeds as follows:

- Compute the average charged hadron multiplicity $< n_{ch} >$ for a given invariant mass (W) using empirical expressions of the ($< n_{ch} >= a_{ch} + b_{ch} * lnW^2$) form. The coefficients a_{ch}, b_{ch}, which depend on the initial state, have been determined by bubble chamber experiments [84].
- Compute the average total hadron multiplicity $< n_{tot} >$ as $< n_{tot} >= 1.5 < n_{ch} >$.
- Using the calculated average hadron multiplicity, generate the actual hadron multiplicity taking into account that the multiplicity dispersion is described by the KNO scaling law, ($< n > P(n) = f(n/ < n >)$) [85]. P(n) is the probability of having n hadrons in the final state given an expected average of $< n >$, and $f(n/ < n >)$ is a universal scaling function. The KNO scaling is parametrized employing the Levy [2] function with an input parameter c_{ch} that depends on the initial state and is treated as a tuning parameter.
- Generate hadrons up to the generated hadron multiplicity taking into account the hadron shower charge conservation and the kinematical constraints. Protons and neutrons are produced in the ratio 2:1 for νp interactions, 1:1 for νn and $\bar{\nu} p$, and 1:2 for $\bar{\nu} n$ interactions. Charged mesons are then created in order to balance charge, and the remaining mesons are generated in neutral pairs. The probabilities for each are 31.33% (π^0, π^0), 62.67% (π^+, π^-) and 6% of strange meson pairs. The probability of producing a strange baryon via associated production is determined from a fit to Λ production data [86, 87, 88, 89]:

$$P_{hyperon} = a_{hyperon} + b_{hyperon} \ \ln W^2 \tag{2}$$

Fig. 1 shows the data/model comparisons of the average charged hadron multiplicity $< n_{ch} >$ as a function of the squared hadronic invariant mass for νp and νn interactions. Fig. 2 shows the data/model comparisons of the negatively charged hadron multiplicity dispersion D_- as a function of the average charged hadron multiplicity $< n_- >$ and the reduced dispersion $D_-/ < n_- >$ as a function of the squared hadronic invariant mass.

The main dynamical feature observed in the study of hadronic systems is that the baryon tends to go backwards in the hadronic center of mass and that the produced hadrons have small transverse momentum relative to the direction of momentum transfer. These features are naturally described in the quark model where the baryon is formed from the diquark remnant, which goes backwards in the center-of-mass, and transverse momentum is generated solely through intrinsic parton motion and gluon radiation. At low invariant masses energy-momentum constraints on the available phase space play a particularly important role. The most pronounced kinematical feature in this region is that one of the produced particles (proton or neutron) is much heavier than the rest (pion and kaons) and exhibits a strong directional anticorrelation with the momentum transfer.

Our strategy, therefore, is to correctly reproduce the final state nucleon momentum, and then perform a phase space decay on the remnant system employing, in addition, a p_T-based rejection scheme designed to reproduce the expected meson transverse momentum distribution. The nucleon momentum is generated using input p_T^2 and $x_F = (p_L^*/p_{Lmax}^*)$ PDFs which are parametrized based on experimental data [90, 91]. Once the baryon momentum is selected the remaining particles undergo a phase space decay. The phase space decay employs a rejection method suggested

[2] The Levy function $Levy(z;c) = 2e^{-c}c^{cz+1}/\Gamma(cz+1)$

in [92], with a rejection factor e^{-A*p_T} for each meson. This causes the transverse momentum distribution of the generated mesons to fall exponentially with increasing p_T^2, controlled by the adjustable parameter A which has a default value of 3.5 GeV^{-1}. Here p_T is the momentum component perpendicular to the current direction. One of the remaining challenges in this model, which will be addressed in the future, is to better describe forward and backward hemisphere multiplicity distributions. The forward/backward multiplicity distributions yield an unphysically rapid transition to the PYTHIA values, a feature not seen in other recent hadronization models [93].

Fig. 3 shows the data/model comparisons of the fragmentation function for positively and negatively charged hadrons. 2-body hadronic systems are a special case: The hadronic system 4-momenta are generated by a simple unweighted phase space.

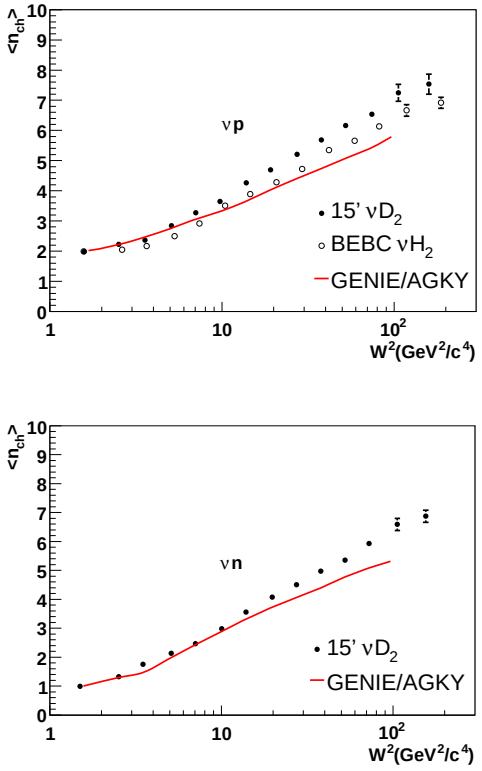

Figure 1: Data/model comparisons of the average charged hadron multiplicity $< n_{ch} >$ shown as a function of the squared hadronic invariant mass for νp and νn interactions. Data are from Refs. [84, 94].

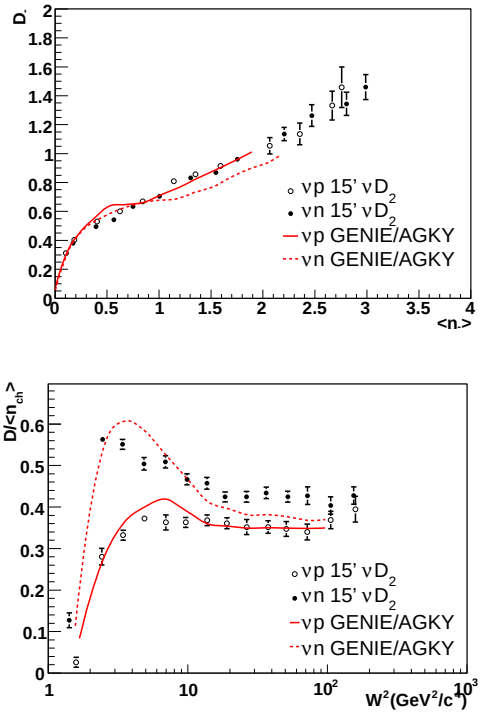

Figure 2: Data/model comparisons of the negatively charged hadron multiplicity dispersion D_- as a function of the average charged hadron multiplicity $< n_- >$ (up) and the reduced dispersion $D_-/ < n_- >$ as a function of the squared hadronic invariant mass (down). Data are from Ref. [84].

3.3.1. Intranuclear rescattering

The hadronization model describes particle production from free targets and is tuned primarily to bubble chamber data on hydrogen and deuterium targets [84, 90, 94, 96, 97, 95, 98, 99]. Hadrons produced in the nuclear environment may rescatter on their way out of the nucleus, and these reinteractions significantly modify the observable distributions in most detectors.

It is also well established that hadrons produced in the nuclear environment do not immediately reinteract with their full cross section. The basic picture is that during the time it takes for quarks to materialize as hadrons, they propagate through the nucleus with a dramatically reduced

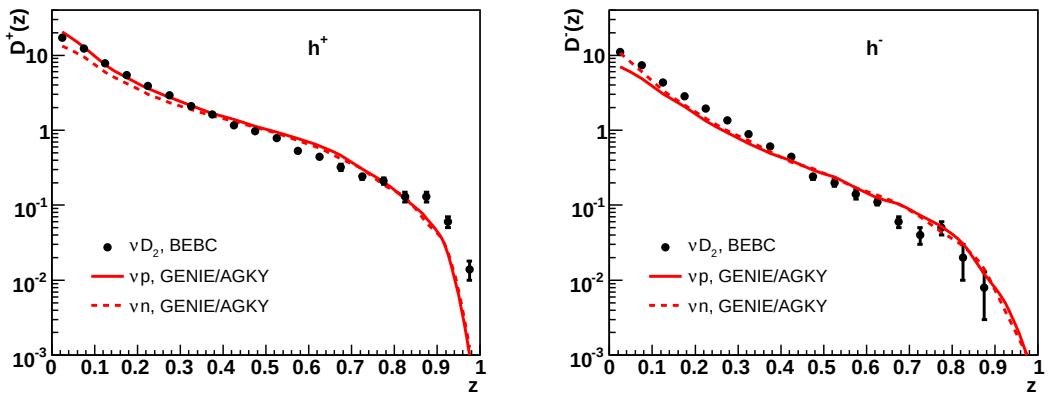

Figure 3: Data/model comparisons of the fragmentation function for positively and negatively charged hadrons. Applied cuts: Squared hadronic invariant mass W^2 above 5 GeV^2/c^4 and squared 4-momentum transfer Q^2 above 1 GeV^2/c^2. Data are from Ref. [95].

interaction probability. This was implemented in GENIE as a simple 'free step' at the start of the intranuclear cascade during which no interactions can occur. The 'free step' comes from a formation time of 0.523 fm/c according to the SKAT model [100].

Intranuclear hadron transport in GENIE is handled by a subpackage called INTRANUKE. INTRANUKE is an intranuclear cascade simulation and has gone through numerous revisions [101] since the original version was developed for use by the Soudan 2 Collaboration [102]. The sensitivity of a particular experiment to intranuclear rescattering depends strongly on the detector technology, the energy range of the neutrinos, and the physics measurement being made. INTRANUKE simulates rescattering of pions and nucleons in the nucleus.

In principle one would like to have a fully realistic nuclear model which accurately describes the full range of processes to model particle production with energies as low as 1 MeV to ensure that the simulations are suitable for any conceivable experiment. Nuclear simulations of this type are themselves highly complex packages and the integration of these packages with GENIE is an area of active work. An alternative approach is to develop a simpler nuclear model, in the context of a particular experiment, and ensure that the relevant physics for that experiment are correctly described. This approach has the advantage of yielding a far simpler code, which is understood by the experimenters. This has particular advantages for the study of systematic errors and the development of ancillary code like reweighting packages.

The current version was optimized for use by the MINOS experiment. For this experiment the task of developing an intranuclear rescattering model is simplified because the detector is composed largely of a single element, iron, and the detector is designed to make a calorimetric energy measurement rather than track individual particles. For the oscillation measurement of MINOS [6] the primary goal is ensuring that the 'missing energy' lost in the nuclear environment is being reliably simulated. The model has applicability to almost all nuclei and a wide range of energies.

To handle a wide energy range for neutrinos, GENIE has defined processes for hadrons up to 1.8 GeV kinetic energy in terms of all the relevant cross sections. For higher hadron energy, the underlying cross sections are assumed to be constant at the 1.8 GeV value. This is a good approximation to the actual values. The code then has a description of all hadrons coming from neutrinos at all relevant energies.

The simulation tracks particles through the nucleus in steps of 0.05 fm. For each particle only one reinteraction is allowed, and the simulation consists of the following steps:

1. **Mean Free Path:** In order to determine if the particle interacts in a particular step the mean free path (λ) is calculated based on the local density of nucleons ($\rho(r)$) [36] and a partial wave analysis of the large body of hadron-nucleon cross sections (σ_{hN}) [103]:

$$\lambda(r, E_h) = \frac{1}{\sigma_{hN}(E_h)\rho(r)} \tag{3}$$

. Here, σ_{hN} is the isospin averaged cross section for the propagating hadron (with energy E_h) interacting with a nucleon and $\rho(r)$ is the matter density of nucleons at the position of the propagating hadron. We use charge densities which are well-measured and are known to be very similar to the matter densities. The code presently tracks pions and nucleons. Isospin relations are used to estimate π^0-nucleon reactions. All nuclei heavier than oxygen are modeled with a Woods-Saxon density distribution and lighter nuclei are modeled with a modified Gaussian distribution.
One difficulty in this approach is that our treatment is using a semiclassical model to describe a quantum mechanical process. This poses particular difficulty in describing elastic scattering which dominates the total cross section at low energy. This wave/particle distinction depends on energy with lower energy hadron-nucleus scattering being more wave-like. To account for this we increase the size of the nuclear density distribution through which the particle is tracked by an amount

$$f\frac{hc}{p}, \tag{4}$$

where f is an adjustable dimensionless parameter set to 0.5 for pions and 1.1 for nucleons in the current default. We use the isospin-averaged total cross sections for pions and nucleons.
2. **Determining the Interaction Type:** Hadron-nucleus interactions occur with different processes and each has an associated cross section - σ_{elas} for elastic scattering (residual nucleus in the ground state), σ_{inel} for inelastic scattering (residual nucleus in excited state, typical response is single nucleon emission), σ_{cex} for single charge exchange (outgoing hadron changes charge, typical response is single nucleon emission) for all hadrons. For pions, emission of 2 or more nucleons with no pion in the final state is called absorption - σ_{abs}; for nucleons, a final state with 3 or more nucleons is called multi-nucleon knockout - σ_{ko}. For low energy nucleons, the knockout mechanism dominates. At higher energies (above about 400 MeV for pions and 600 MeV for nucleons), the probability of pion production (σ_{piprod}) becomes important. The total cross section (σ_{tot}) is the sum of all component cross sections and the reaction cross section (σ_{reac}) is the sum of cross section for all inelastic reactions,

$$\sigma_{reac} = \sigma_{cex} + \sigma_{inel} + \sigma_{abs} + \sigma_{piprod} = \sigma_{tot} - \sigma_{elas}. \tag{5}$$

This equation is specific to pions; σ_{abs} is replaced by σ_{ko} for nucleons. Once it has been determined that a hadron reinteracts in the nucleus, the type of the interaction is determined based on the measured cross sections for the above listed processes. Cross sections for kinetic energy less than 1 GeV are used and they are assumed constant above 1 GeV. Where data is sparse, cross section estimates are taken from calculations of the CEM03 group [104]. Since only 1 reinteraction is allowed, the effect of additional reactions with the rest of the nucleus must be included here. The distribution of final states is optimized for iron, but valid for all targets. Figure 4 shows the default INTRANUKE compared with π^+ data for σ_{tot} and σ_{reac} for iron and carbon. These are examples of many successful comparisons.
3. **Final State Products:** Once the interaction type has been determined, the four-vectors of final state particles need to be generated. Where possible these distributions are parametrized from data or from the output of more sophisticated nuclear models [105]. Simple models are used for elastic scattering angular distributions. The quasielastic reaction mechanism is known to dominate the final state for inelastic or charge exchange processes. Fermi motion and binding energy are used to get a good description of the kinematics. Whenever 3 or more particles are emitted, distributions are according to phase space. For MINOS, the most important issue is missing energy generated by inelastic and absorption processes.

Very low energy hadrons and nuclear recoils are not seen, so simplifications can be made. All states where more than 5 nucleons are emitted are treated as though 5 nucleons (3 protons and 2 neutrons) were emitted. These restrictions are relaxed in the most recent code versions to better match what is seen in data.

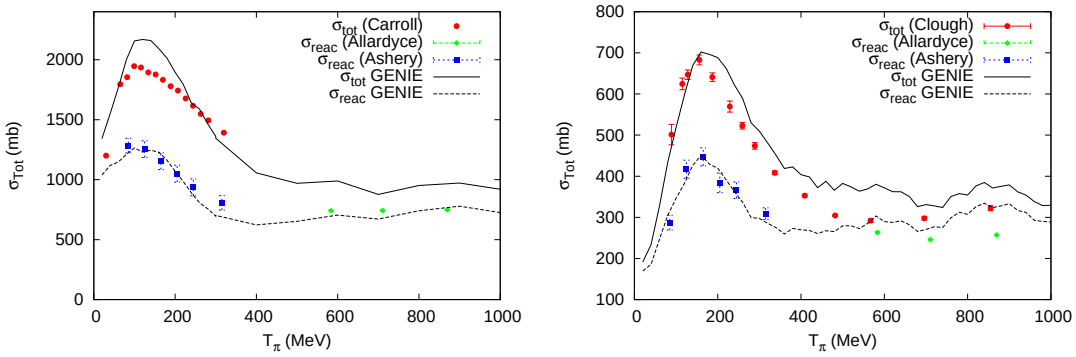

Figure 4: π^+ total and reaction cross sections for iron (left) and carbon (right). Data are from Refs. [106, 107, 108, 109].

3.4. Physics Model Tuning

The full range of models involve more than one hundred adjustable parameters, the complete set of which are given in the Physics and User Manual [110]. Only the most important in the construction of the physics models will be discussed here. Electroweak parameters and CKM matrix elements use the values of the Particle Data Group [49].

As mentioned previously, the quasi-elastic, resonance production, and DIS models employ form factors, axial vector masses, and other parameters which have been determined by others in their global fits [43, 47]. In order to check the overall consistency of our model, and to verify that we have correctly implemented the DIS model, predictions are compared to electron scattering inclusive data [111, 112] and neutrino structure function data [113]. The current default values for transition region parameters are W_{cut}=1.7 GeV/c^2, $R_2(\nu p) = R_2(\overline{\nu} n)$=0.1, $R_2(\nu n) = R_2(\overline{\nu} p)$=0.3, and R_m=1.0 for all $m > 2$ reactions. These are determined from fits to inclusive and exclusive (one and two-pion) production neutrino interaction channels. For these comparisons we rely heavily on online compilations of neutrino data [114] and related fitting tools [115] that allow one to include some correlated systematic errors (such as arising from flux uncertainties). The GENIE default cross section for $\nu\mu$ charged current scattering from an isoscalar target, together with the estimated uncertainty on the total cross section, as evaluated in [116] are shown in Fig. 5.

The tuning of the hadronization model is accomplished using data from the BEBC [94], FNAL [117], and SKAT [118] bubble chamber experiments, and is described in more detail elsewhere [119]. Multiplicity measurements include averaged charged and neutral particle (π^0) multiplicities, forward and backward hemisphere average multiplicities and correlations,topological cross sections of charged particles, and neutral - charged pion multiplicity correlations.Hadronic system measurements include fragmentation functions (z distributions), x_F distributions, p_T^2 (transverse momentum squared) distributions, and $x_F - \langle p_T^2 \rangle$ correlations ("seagull" plots) [120]. Averaged charged particle multiplicity and dispersion parameters are taken from published values [84], as well as our own fits [119]. Baryon 4-momentum distributions are determined from fits to experimental data [90, 91]. The settings for PYTHIA parameters are taken as the non-default values tuned for the NUX [18] generator, a high energy generator used by the NOMAD [121] experiment.

The intranuclear rescattering model has been tested and tuned based on comparisons to hadron-nucleus data. Hadron-nucleus cross sections are calculated by 'MC experiments' where a nucleus

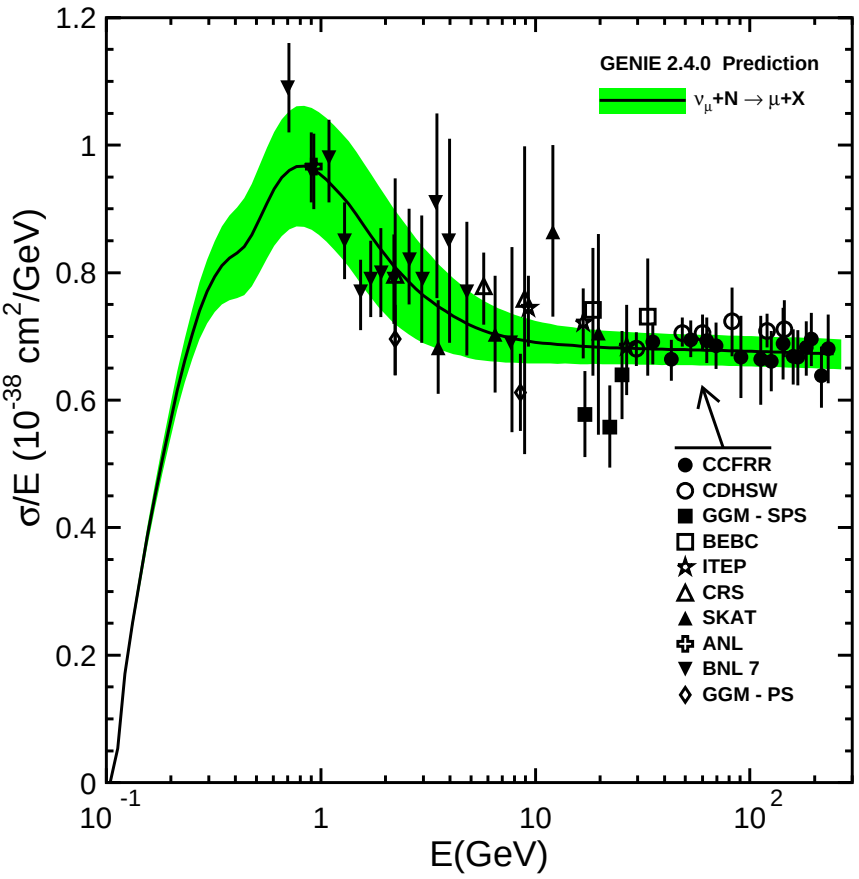

Figure 5: GENIE default cross section for $\nu\mu$ charged current scattering from an isoscalar target. The shaded band indicates the estimated uncertainty on the free nucleon cross section. Data are from [62] (CCFRR), [63] (CDHSW), [64] (GGM-SPS), [65, 66] (BEBC), [67] (ITEP), [68] (CRS, SKAT), [69] (ANL), [70] (BNL) and [71] (GGM-PS)

is being illuminated by a uniform hadron beam with transverse radius larger than the nucleus size. Figure 4 shows the comparison between INTRANUKE and data for π^+-C and π^+-Fe total and reaction cross sections. Extrapolations to higher energy are required in many cases as data for only σ_{reac} is available. CEM03 [105] results with appropriate rescaling to match data at lower energies are often used. Although the model is tuned to hadron scattering on iron, the simplicity of the Fermi Gas model and the $A^{(2/3)}$ scaling of the cross sections allow the model to be applied to nearly all nuclei encountered in the simulation as well.

Validation of the intranuclear rescattering model using neutrino data has also been performed. This revisits the analysis of Reference [102] which compares ANL neutrino scattering data on deuterium [81] to BEBC neutrino-neon data [122], where each have been rescaled to an atmospheric neutrino flux. By comparing neon and deuterium final state topology fractions the rates for pion absorption and charge exchange in neutrino-neon interactions can be determined. INTRANUKE reproduces the measured final state topology fractions with an overall χ^2 of 16.0 for 12 degrees of freedom. The rates of pion absorption and charge exchange produced by INTRANUKE in this comparison are $18.3 \pm 0.5\%$ and $2.9 \pm 0.2\%$ respectively, in good agreement with the measured values of $22 \pm 5\%$ and $10 \pm 8\%$.

In evaluating the uncertainty in the intranuclear rescattering model, several sources of uncertainty

were taken into account for their effect on the MINOS determination of the hadronic energy scale [123]. These include the experimental uncertainty on the external data that serves as the input to the model, as well as on some of the key theoretical assumptions in the model, in particular the modeling of pion absorption reactions and the treatment of low energy pion scattering [123].

4. Recent Developments

Recent focus on development for GENIE has been on the nuclear structure and final state interaction codes. The purpose is to make the code well-tuned to the needs of the upcoming generation of accelerator experiments including T2K [7], NoVA [8], Minerva [9], MicroBooNE [10] and ArgoNEUT [11]. We are also looking to push the validity range of GENIE down to the MeV scale making it applicable for the study of neutrinos from reactors, supernova and SNS [124].

The Fermi Gas nuclear model has been shown to be wrong in detail through interpretation of electron scattering experiments [125]. Nucleon-nucleon correlations are important in kinematic regions where the impulse approximation is unlikely to apply. The spectral function [126, 127] has become a useful model to represent the effects of a many-body model. Developmental versions of GENIE now contain spectral functions of Benhar for carbon, iron, and lead. Code for calculating (e, e') differential cross sections is now in place.

A true internuclear cascade model has also been in development [101]. It tracks pions and nucleons through multiple reactions in the same nucleus in which the neutrino was absorbed. Free hadron-nucleon cross sections are used with the struck nucleon has momentum and binding energy according to the Fermi Gas model. Interactions for protons, neutrons, and pions are presently modeled. The hadrons are on-shell between scatterings; the reactions are governed by the same mean free path and Fermi gas models as for the existing model. Thus, there is no limit on the number of particles that can be tracked. A simple model for compound nuclear processes is included to properly account for effects in hadron-nucleus data.

5. Software Design Overview

In this section we will describe the software design of the GENIE package. We begin by discussing the software requirements and use cases. The software framework is presented together with key classes such as *particles, events,* and *interactions*. The hierarchical delegation of responsibility during event generation to *driver, thread,* and *module* classes is described.

5.1. Requirements

The process of requirements capture and software design was carried out over a period of several months in 2004 and 2005. During this phase there was extensive discussion within the MINOS collaboration as well as in conjunction with the NuINT[3] Conference series. Through these meetings we received input from many users of these packages as well as with several experts who had designed, developed, and supported previous (Fortran-based) procedural codes.

Through these discussions the need for a package of this type for future neutrino experiments became very clear. The procedural programs had often been in use for several decades and had expanded greatly beyond their initial scope. They had often reached a critical mass where further modifications were deemed extremely difficult because of the overall fragility of the architecture. In addition there were often strong couplings between aspects of the simulation that made incremental improvements in a particular area difficult. Another commonly voiced concern was the

[3]Neutrino Interactions in the Few-GeV Energy Range

lack of documentation, particularly regarding the ways in which the models and their parameters were tuned to data.

These discussions served to illuminate several typical use cases for neutrino event generators and related tools:

- For event generation in conjunction with a full detector simulation.
- For event generation in fast simulations, either 4-vector only or using a parametrized model of the detector response.
- To provide a library of cross section values for interaction rate calculations.
- As a source of information about the underlying models and their uncertainties.
- As a primary tool in the evaluation of systematic uncertainties.

These use case evaluations and discussions led to the establishment of a set of requirements for the overall architecture as well as requirements for the documentation and ancillary tools:

- Decouple physics models as much as possible from the framework code.
- Lower the barrier to entry for physics model developers (particularly theorists).
- Provide a well-tested, well-defined default configuration which provides a benchmark for all users who are primarily interested in using the package in black box-mode.
- Incorporate up-to-date theoretical and experimental work and provide a flexible framework so that it can be maintained.
- Incorporate state-of-the-art software engineering methodologies to support these goals through an object-oriented design.
- Leverage the developments in other areas of HEP software development (in particular ROOT class libraries).
- Provide the external data used to tune the package as part of the overall distribution.
- Provide clear documentation about how the models are tested, tuned, and validated.
- Provide a set of tools to facilitate the tuning of model parameters and for an independent evaluation of how well models describe existing data.

These requirements were met in the August 2007 first release of the GENIE package. The implementation of the physics models was cross-checked through an exhaustive set of comparisons [128] with one of the existing procedural codes [13]. GENIE can be configured to be identical in its physics content to the final version (v3.5.5) of the NEUGEN3 package, one of the legacy procedural codes which had been used by numerous experiments for over more than a decade.

5.2. Core Framework

The key requirement of the GENIE Core Framework is to transparently decouple the high-level code focusing on physics simulations from the low-level structures involved primarily with memory management and configuration.

The framework was developed and reviewed primarily within the MINOS experiment and, inevitably, has been influenced by the MINOS offline software design [129]. In developing the GENIE framework we recycled, adapted and extended key features of the MINOS offline framework.

We drew heavily from the accumulated software engineering experience encapsulated within popular software design patterns, including the Visitor, Chain of Responsibility, Factory, Strategy and Singleton [130]. The Core Framework is not specific to the subject matter domain of GENIE and could be adapted and reused in other scientific computing applications. The GENIE Event Generation Framework, to be discussed later, is a subject matter-specific layer built on top of the Core Framework.

The framework concerns itself mainly with the properties, instantiation and memory management of software abstractions called *Algorithms*, and specifies the interfaces that underpin the interactions between the numerous concrete *Algorithm* realizations. The *Algorithm* is a key Core Framework abstraction. The notion of an 'algorithm' in an object-oriented system requires further clarification as it does not correspond to its more familiar notion in the context of procedural software systems. In the Core Framework the *Algorithm* encapsulates the common behavior of all algorithmic objects. It is an abstract base class which defines exactly how algorithmic objects are to be initialized and configured, how they are to look up their configuration, how they are to be identified, and how they report their status. These are common, largely operational features that characterize a very heterogeneous collection of algorithms such as cross section models, hadronization models, particle decayers, form factor and structure function models, event generation modules and threads and other types of algorithms that can be found within GENIE. The fact that such a common behavior is imposed upon all algorithmic objects allows us to build a central, external, XML [131] - based algorithm configuration system that contributes significantly to the flexibility and extensibility of GENIE. The kind of computation to be performed, the usual identifying feature of an algorithm in a procedural system, is a secondary characteristic at this level of abstraction.

At the next level up from the *Algorithm* root of an algorithm inheritance tree, we find a standardized interface which defines how to invoke each specialized type of calculation and retrieve its results. Numerous such specialized algorithm interfaces exist within GENIE. Examples include the *GFluxI* interface for concrete flux drivers, the *XSecAlgorithmI* interface for concrete cross section models, and the *EventRecordVisitorI* interface for for concrete event generation steps. Invoking all algorithms through such standardized interfaces guarantees scalability and ensures the seamless integration of new concrete implementations.

The algorithmic objects are stateless and their behavior is fully externally configured. The algorithm configurations are stored in XML files. Typically, there is a single XML configuration file per algorithm. Each file may contain multiple configuration sets for that algorithm with each configuration set being uniquely identified by a name. The algorithm configuration variables can be of many different types (including booleans, integers, real numbers, strings, ROOT 1-D or 2-D histograms, ROOT n-tuples/trees or other GENIE algorithms with their own configurations). Each configuration variable, in a given set, is uniquely identified by a name. During the initialization phase, all XML configuration files are parsed and each named configuration set is stored at a type-safe 'value' $\rightarrow$ 'type' associative container called the *Registry* . All *Registry* objects instantiated in initialization phase are stored in a shared pool called the *AlgConfigPool*. A unique name is being used to identify each *Registry* in that pool. The name is constructed by the name of the configuration set, the name of the algorithm the configuration is intended for and the namespace that the algorithm lives in, as 'namespace::algorithm-name/configuration-name'. At run-time each algorithmic object can look up its configuration set by accessing the corresponding *Registry* object.

One feature of the GENIE configuration system is especially worth noting. Algorithm configuration sets may include other algorithms (with their own configurations, which in turn may contain more algorithms). GENIE's extensibility and flexibility is largely due to this feature in conjunction with the standardization of the algorithm interfaces. In the actual GENIE code one only needs define a call sequence between abstract algorithm-types such as, for example, that an algorithm-type specialized in generating scattering kinematics, invokes another algorithm-type specialized in cross section calculations which, in its turn, should invoke another algorithm type

specialized in form factor calculations. Once that call sequence has been defined in the code, many concrete realizations may come into being purely at the configuration level by specifying the names of the concrete algorithms and the names of their configuration sets.

Typically, pre-configured instances of GENIE algorithms are accessed through an algorithm factory [130] which is responsible for instantiating each algorithm (upon request) and allowing it to look up its configuration. The factory typically owns and manages the list of all instantiated concrete algorithms. Since algorithms are stateless objects, further requests for an instantiated concrete algorithm results in the previously instantiated algorithm being returned rather than a new one being created.

By default all instantiated concrete algorithms and configurations are stored within shared pools designed as singletons [130]. As these are shared pools, modifications have global effects. For example, modifying a low-level algorithm configuration modifies all call sequences that include that algorithm. This is desirable in most contexts, such as for example for the consistent propagation of physics parameter changes throughout GENIE. There are certain situations, however, such as fitting or event reweighting applications, where this may be not be a desirable feature. The GENIE Core Framework allows algorithms to clone and assume ownership of the entire sequence of sub-algorithms they depend upon, along with each sub-algorithm's configuration registries. That cloned call-sequence of algorithmic objects is stored in a local rather than a shared pool. In this way, concrete top-level algorithms behave as self-contained capsules and can be re-configured in isolation without affecting other GENIE components.

5.3. Event generation framework: Particles, Events and Interactions

In this section three key framework classes, the *GHepParticle*, *GHepRecord*, and the *Interaction* classes, are described.

GENIE is using the natural system of units ($\hbar = c = 1$) so almost every simulated quantity is expressed in powers of [GeV]. Exceptions are the event vertex in the detector coordinate system (in SI units) and particle positions in the hit nucleus coordinate system (in fm). Different units may be employed when native GENIE event descriptions are converted to experiment-specific formats in accordance with the format specification.

5.3.1. Particles

The basic output unit of the event generation process is a 'particle'. This is a term used to describe both particles and nuclei appearing in the initial, intermediate or final state, as well as generator-specific pseudo-particles used for facilitating book-keeping of the generator actions.

Each such 'particle' generated by GENIE is an instance of the *GHepParticle* class. These objects contain information with particle-scope including: particle ID and status codes, PDG mass, charge, name, indices of mother and daughter particles marking possible associations with other particles in the same event, 4-momentum, 4-position in the target nucleus coordinate system, polarization vector, and other properties. The *GHepParticle* class includes methods for setting and querying these properties.

GENIE has adopted the standard PDG particle codes [49]. For ions it has adopted a PDG extension, using the 10-digit code 10LZZZAAAI where L is the number of strange quarks ZZZ is the total charge, AAA is the total baryon number, and I is the isomer number (I=0 corresponds to the ground state). GENIE-specific pseudo-particles have PDG code >= 2000000000 and can convey important information about the underlying physics model. Pseudo-particles generated by other specialized programs that may be called by GENIE (such as PYTHIA-6) are allowed to retain the codes specified by that program.

GENIE obtains particle data (including particle names, codes, masses, widths, decay channels and more) using the ROOT's *TDatabasePDG*. This particle data-base manager object is initialized

with the constants used in PYTHIA-6. The data-base has been augmented by the GENIE authors to include baryon resonances, nuclei and GENIE-specific pseudo-particles. Details are given in Ref. [110].

GENIE marks each particle with a status code. This signifies the position of a particle in a time-ordering of the event and helps navigation within the event record. Most generated particles are marked as one of the following:

- 'initial state', typically the first two particles of the event record corresponding to the incoming neutrino and the nuclear target.
- 'nucleon target', corresponding to the hit nucleon (if any) within the nuclear target.
- 'intermediate state', typically referring to the remnant nucleus, fragmentation intermediates such as quarks, diquarks, or intermediate pseudo-particles.
- 'hadron in the nucleus', referring to a particle of the primary hadronic system, defined as the particles emerging from the primary interaction vertex before any possible re-interactions in the nucleus.
- 'decayed state', such as for example unstable particles that have been decayed.
- 'stable final state' for the relatively long-lived particles emerging from the nuclear targets.

All particles generated by GENIE during the simulation of a single neutrino interaction are stored in a dynamic container representing an 'event'.

5.3.2. Events

Events generated by GENIE are stored in a custom, STDHEP-like event record called a *GHEP* record. Each *GHEP* event record, an instance of the *GHepRecord* class, is a ROOT TClonesArray container of *GHepParticle* objects representing individual particles.

Other than being a container for the generated particles, the event record holds additional information with event-, rather than particle-, scope such as the cross sections for the selected event, the differential cross section for the selected event kinematics, the event weight, a series of customizable event flags, and interaction summary information (described in the next section).

Additionally, the event record includes a host of methods for querying / setting event properties including many methods that allow querying for specific particles within the event. Examples include methods to return the target nucleus, the final state primary lepton, or a list of all stable descendants of any intermediate particle.

The event record features a 'spontaneous re-arrangement' feature which maintains the compactness of the daughter lists at any given time. This is necessary for the correct interpretation of the stored particle associations as the daughter indices correspond to a contiguous range. The particle mother and daughter indices for all particles in the event record are automatically updated as a result of any such spontaneous particle rearrangement.

The event generation itself is built around the *GHEP* event record using the Visitor design pattern [130]. The interaction between the *GHEP* event record and the event generation code will be outlined in the following sections.

The *GHEP* structure is highly compatible with the event structures used in most HEP generators. That allows us to call other generators, such as PYTHIA-6, as part of an event generation chain and convert / append their output into the current *GHEP* event. Additionally the *GHEP* events can be converted to many other formats for facilitating the GENIE interface with experiment-specific offline software systems and cross-generator comparisons.

5.3.3. Interactions

The *GHEP* record represents the most complete description of a generated event. Certain external heavy-weight applications such as specialized event reweighting schemes or realistic detector simulation chains using the generator as the physics front-end require all of the detailed particle-level information. However, many of the actual physics models employed by the generator, such as cross section, form factor, or structure function models, require a much smaller subset of information about the event.

An event description based on simple summary information, typically including a description of the initial state, the process type and the scattering kinematics, is sufficient for driving the algorithmic objects implementing these physics models. In the interest of decoupling the physics models from event generation and the particle-level event description, GENIE uses an *Interaction* object to store summary event information. Whenever possible, algorithmic objects implementing physics models accept a single *Interaction* object as their sole source of information about an event. This enables the use of these models both within the event generation framework but also within a host of external applications such as model validation frameworks, event re-weighting tools and user physics analysis code.

An *Interaction* object is an aggregate, hierarchical structure, containing many specialized objects holding information for the initial state (*InitialState* object), the event kinematics (*Kinematics* object), the process type (*ProcessInfo* object) and potential additional information for tagging exclusive channels (*XclsTag* object). Instantiating *Interaction* objects for driving physics models is streamlined using the 'named constructor' C++ idiom. They can be serialized into unique string codes which, within the GENIE framework, play the role of the 'reaction codes' of the old procedural systems. These string codes are used extensively for mapping information to interaction types. Two examples include mapping interaction types to pre-computed cross section splines or mapping interaction types to specialized event generation code. Each generated event has an *Interaction* summary object already attached to it.

5.4. Event generation processing units: Modules, Threads and Drivers

On an operational level the responsibility for generating events is shared between event generation *drivers, threads* and *modules.* Tasks are delegated from event generation drivers to threads, and from threads to modules. Event generation drivers can include multiple threads, and threads can include multiple modules. Event generation drivers are responsible for generating events for a particular user-defined situation. These can be as simple as monoenergetic neutrinos interacting off a single target, to complex situations involving the output of realistic beam-line simulations and full detector geometry descriptions. Threads are responsible for generating the physics of particular classes of events, for instance charged-current quasielastic. Modules carry out a single step in that generation process.

5.4.1. Event generation modules

An event generation module is a key event generation abstraction. Each event generation module encapsulates a well defined event generator operation which, in physics terms, can be any of a very diverse set of actions. Examples include selecting the scattering kinematics, generating the final state primary lepton or the primary hadronic system, transporting hadrons within the target nucleus, and decaying unstable particles.

Operationally, event generation can be seen as a series of well-defined processing steps operating on the *GHEP* event record. The act of operating on the event record defines an interface that is encapsulated by the *EventRecordVisitorI* abstract class. As it is indicated by the interface name, the Visitor design pattern is being employed [130]. Concrete event generation modules, implementing the *EventRecordVisitorI* interface, 'visit' the event record. The event record then invokes each attached module and is modified as a result.

Due to the diversity of the event processing operations that must be considered by GENIE, we formed the event generation module abstraction focusing on the common operational aspect of (potentially) modifying the event record. This represents the most generic way of thinking about event generation and guarantees that any future physics addition, especially ones not envisioned at this stage of the generation evolution, can be trivially embedded into the existing framework. Treating the event generator modules uniformly and standardizing on the event generation module interface allows us to build a flexible and extensible system where modules can be dynamically plugged in/out of the event generation or interchanged. Examples can further clarify the utility of this abstraction: a module handling a set of particle decays can be unplugged to inhibit those decays, or a module handling intra-nuclear hadron transport may be swapped with another module performing the same operation using a different physics strategy.

Whenever possible event generation modules are written in a generic way, containing code implementing just the neutrino event generation mechanics. The actual physics model itself is specified in the generation module configuration. This decoupling of mechanics from models greatly simplifies code development, transparency, and physics validation, simplifying the overall structure and reducing the amount of code that needs to be actively developed and scrutinized between successive releases. An example will clarify this factorization: The module selecting the kinematics for deep-inelastic neutrino-nucleon interactions does not contain the actual code for the deep-inelastic differential cross section. It merely contains code to calculate the allowed kinematical phase space for the process, select a point in that phase space using a Monte Carlo acceptance / rejection method, and update the *GHEP* record accordingly. The actual differential cross section model used during the Monte Carlo selection is an external physics model invoked by the event generation module. The module itself can be recycled many times by instructing it to call a different cross section model each time. As a result of that factorization, multiple call sequences can be defined purely at the configuration level without code duplication.

5.4.2. Event generation threads

An *event generation thread* is an ordered sequence of processing steps, encapsulated by event generation modules, that can be applied to an empty *GHEP* event record to completely generate some class of physics events. This process defines an interface that is encapsulated by the *EventGeneratorI* abstract class. Within the GENIE event generation framework the structures containing a comprehensive set of instructions for generating a class of physics events are concrete *EventGeneratorI* objects.

GENIE defines a comprehensive set of event generation threads responsible for generating event types at the level of fundamental interactions. The complete set of these event generation threads comprises GENIE's full 'physics content' for event generation. As an event generation thread can generate a single class of events only, there are usually multiple threads in use.

The class of physics events generated by a thread can have an arbitrary granularity, from a single interaction corresponding to a particular process type with a given final state to very broad event categories. Each thread contains an *InteractionList* object, a container containing a list of the *Interaction* objects the thread can generate. The *InteractionList* plays a crucial role in identifying the responsibilities of each thread within the GENIE framework. Once an event type to be generated has been selected, a corresponding *Interaction* object is instantiated. Following the Chain of Responsibility design pattern [130], GENIE attempts to match the *Interaction* object with an element of the *InteractionList* containers for all active threads. The first thread found that is able to handle that event type is handed the responsibility to generate the event.

Additionally, event generation threads include an instance of the cross section algorithm that can be used for selecting the event kinematics or for computing the probability for a particular neutrino to interact. This is another example of separating mechanics from models and serves to greatly simplify the dynamic mapping between event types and cross section models.

Once a list of threads has been loaded into the generator, many high-level event generation op-

erations became trivial. Compiling the list of all event types that can be generated by GENIE in its current configuration simply involves looping over the active threads and adding the corresponding *InteractionList* objects. Selecting an event type to be generated from that master list involves looping over its *Interaction* objects and, for each element, identifying the responsible thread, requesting its corresponding cross section model and invoking it by passing the *Interaction* object as argument. Once an event type has been selecting generating the event simply involves looking up the responsible thread and delegating responsibility to it.

During event generation an invoked thread maintains a modification history of the event record. If a tried event generation path leads to a dead-end, the current event generation module throws an exception and aborts. The event generation thread catches that exception and, depending on information stored at it, may rerun the event using a snapshot of the event record taken N steps back, in the hopes of taking an alternative path and avoiding the encountered dead-end. If a configurable maximum number of exceptions is caught, or if any thrown exception specifies explicitly that generation of the current event is to be aborted altogether, the thread sets the appropriate error flags and makes sure that the remaining processing steps are skipped. The user, via options set in the event generation driver, may choose to keep certain types of these events so as to examine their type and frequency, though the default behavior of GENIE is to discard these events and only write out physical, fully formed events. Error handling within each active thread greatly adds to the robustness and fault-tolerance of the package, which is especially valued in large-scale, CPU-intensive, experiment-specific Monte Carlo production runs involving hundreds of CPU cores over many weeks.

Advanced users can modify the default event generation threads by removing / adding event generation modules, or they can define their own uniquely named threads for handling new processes or handling existing processes in new ways.

5.4.3. Event generation driver classes

GENIE provides two event generation driver classes. These drivers collect the user inputs, instantiate and configure all required event generation components, and oversee communications between these components, the computing environment, and the user.

The two driver classes support two different types of functionality:

- Instances of the *GEVGDriver* class can handle event generation for a given initial state corresponding to an arbitrary neutrino / target pair.
- Instances of the *GMCJDriver* class can be used for more complicated simulations involving arbitrarily complex, realistic beam flux simulations and detector geometry descriptions. This driver object concerns itself mostly with driving the flux and detector geometry navigation drivers and integrating those with the GENIE event generation framework. It represents a significantly more complex and CPU-intensive event generation case but relies entirely on a host of *GEVGDriver* instantiations, one for each possible initial state in that Monte Carlo job, in order to obtain neutrino interaction physics modeling capabilities and generate event kinematics.

6. GENIE Event Generation Applications and Utilities

GENIE is being used by a host of precision-era neutrino experiments and provides off-the-shelf components for generating neutrino interactions under the most realistic assumptions. The event generation driver classes described in Section 5 are encapsulated within driver applications which present the user with a command-line or graphical interface, instantiate and configure those driver

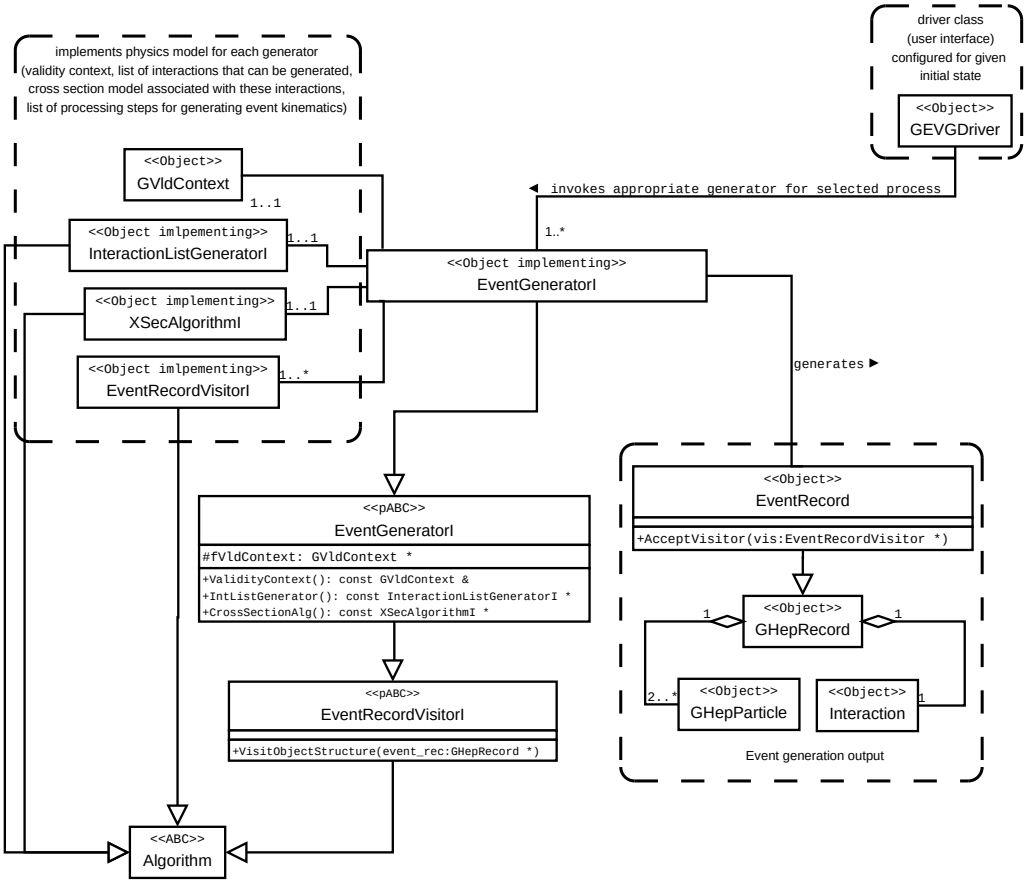

Figure 6: A UML diagram depicting the GENIE event generation framework. See text for details.

classes, call the event generation methods to generate the requested number of events, and push those events through a persistency manager.

In experiment-specific GENIE-based event generation drivers utilizing the *GMCJDriver* one can integrate the GENIE neutrino interaction modeling with detailed flux and detector geometry descriptions. This is a non-trivial operational capability that older procedural neutrino generators typically lacked, requiring significant development effort from experiments. The flux descriptions are typically derived from experiment-specific beam-line simulations while the detector geometry descriptions are typically derived from engineering drawings mapped into the GEANT4 [2], ROOT [12] or GDML [132] geometry description languages. Obviously, flux and detector geometry descriptions can take many forms, driven by experiment-specific choices. GENIE standardizes the geometry and flux driver interfaces by defining the operations that GENIE needs to perform on the geometry and flux descriptions and the essential flux and geometry information needed for the generation of events. Concrete implementations of the interfaces allow experiments to extend GENIE's event generation capabilities and make it possible to seamlessly integrate new geometry descriptions and beam fluxes into user applications.

In this section we will describe in some detail the flux and geometry interfaces. We will briefly describe applications built from these drivers as well as GENIE utilities to evaluate and display cross section information, make comparisons to external data, and facilitate model tuning.

6.1. Neutrino flux drivers

In GENIE every concrete flux driver implements the *GFluxI* interface. The interface defines what neutrino flux information is needed by the event generation drivers and how that information is to be obtained. Each concrete flux driver includes methods to:

- Declare the neutrino flavors that can generate events. This information is used for initialization purposes, in order to construct a list of all possible initial states in a given event generation run.
- Declare the maximum energy. Again this information is used for initialization purposes, in order to calculate the maximum possible interaction probability in a given event generation run. Since neutrino interaction probabilities are tiny, GENIE scales all interaction probabilities in a particular event generation run so that the maximum possible interaction probability is 1. That maximum interaction probability corresponds to the total interaction probability (summed over nuclear targets and process types) for a maximum energy neutrino following a trajectory that maximizes the density-weighted path-lengths for each nuclear target in the geometry. GENIE adjusts the MC run normalization accordingly to account for this probability renormalization.
- Generate a flux neutrino and specify its pdg code, its weight (if any), its 4-momentum and 4-position. The 4-position is given in the detector coordinate system (as specified by the input geometry). Each such flux neutrino is propagated towards the detector geometry but is not required to cross any detector volume. GENIE will take that neutrino through the geometry, calculate density-weighted path-lengths for all nuclear targets in the geometry, calculate the corresponding probabilty for scattering off each nuclear target and decide whether that flux neutrino should interact. If it interacts, an appropriate *GEVGDriver* will be invoked to generate the event.
- Notify that no more flux neutrinos can be thrown. Flux drivers that use the output of beamline simulations, so-called 'flux files', are configured to recycle these flux files multiple times in a given run since most neutrino flux entries do not produce an interaction. The flag allows GENIE to properly terminate the event generation run once this limit is reached, irrespective of the accumulated number of events, protons on target, or other metric of exposure.

The above correspond to the common set of operations/information that GENIE expects to be able to perform/extract from all concrete flux drivers. Specialized drivers may define additional information that can be utilized in experiment-specific applications. One typical example of this is to 'pass-through' information about the flux neutrino parents placed in the flux files by the beamline simulation, such as the parent meson PDG code, its 4-momentum, and its 4-position at the production and decay points.

At the time of writing this article, GENIE already includes a host of concrete flux drivers allowing GENIE to be used in many realistic, experiment-specific situations. More specifically, it includes an interface to the JPARC neutrino beam simulation [133] used by Super-Kamiokande [134], nd280 [7], and INGRID [7] and an interface to the NuMI beam simulation [135] used by MINOS [136], NOvA [8], MINERvA [9], MicroBooNE [10] and ArgoNEUT [11]. It includes drivers for the BGLRS [137] and the FLUKA [138] atmospheric fluxes. Moreover, it includes a generic flux driver, describing a cylindrical neutrino flux of arbitrary 3-D direction and radius, with a configurable radial dependence, which can be used for describing a flux containing a number of different neutrino species whose (relatively normalized) energy spectra are described as 1-D histograms. GENIE, obviously, also includes the trivial case of a monoenergetic flux typically employed in physics benchmarking calculations.

Concluding this section it is worth re-emphasizing the fact that new concrete flux drivers (describing the neutrino flux from other beam-lines) can be easily developed and they can be effortlessly and seamlessly integrated within the GENIE event generation framework.

6.2. Geometry drivers

In GENIE every concrete geometry driver implements the *GeomAnalyzerI* interface. The interface specifies what information about the input geometry is relevant to the event generation and how that information is to be obtained. Each concrete geometry driver implements methods to:

- Declare the list of target nuclei that can be found in the geometry. This information is used for initialization purposes, in order to construct a list of all possible initial states in a given event generation run.
- Compute the maximum density-weighted path-lengths for each nuclear target in the geometry. Again, this is information used for initialization purposes. The computed 'worst-case' trajectory is used to calculate the maximum possible interaction probability in a particular event generation run. This maximum interaction probability is used internally to normalize all computed interaction probabilities in that run.
- Compute density-weighted path-lengths for all nuclear targets, for a trajectory of given 4-momentum and starting 4-position. This allows GENIE to calculate probabilities for each flux neutrino to be scattered off every nuclear target along its path through the detector geometry.
- Generate a vertex along a trajectory of given 4-momentum and starting 4-position on a volume containing a given nuclear target. This allows GENIE to place a neutrino interaction vertex within the detector geometry once an interaction of a flux neutrino off a selected nuclear target has been generated.

GENIE currently contains a concrete geometry driver able to handle the ROOT-based detector geometry descriptions typically used by most neutrino experiments. Detector geometry descriptions based on GEANT or GDML can be converted into ROOT-based descriptions and used by the same GENIE geometry driver as well. GENIE also includes a driver for more trivial geometry descriptions corresponding to a single nuclear target or a target mix (a set of nuclear targets, each with its corresponding weight fraction) at a fixed position. This simpler geometry driver may be used in simulating fixed initial states for benchmarking calculations or in experimental situations where a relatively uniform detector is being illuminated by a spatially uniform neutrino beam. An example of the latter would be a detector placed far enough from the beam-line instrumentation so as to see a point-like neutrino source.

Again it is worth re-emphasizing that any new detector geometry description can be seamlessly integrated with the GENIE event generation framework by means of developing an appropriate GENIE geometry driver.

6.3. Event generation outputs

The generated events are stored in the ROOT file format. The typical output of an event generation run is a single ROOT file which contains an event tree with a single branch and a single leaf per event containing the generated *GHEP* record. User-defined branches to write out experiment-specific information may be added to that tree with the user-defined information 'linked' to the corresponding generated neutrino event. The output ROOT file contains directories storing all GENIE configuration information for the MC run and a snapshot of the running environment for later reference. GENIE provides a persistency manager object which can be employed within the event generation driver applications to write out the event tree.

6.4. Event generation applications

GENIE is distributed with many event generation applications and utilities, many of which are straightforward wrappers for the components described above. Users interact with these applications through simple command-line interfaces, user-created XML configuration files, and environmental variables.

The *gevgen* application can be used in simulating given initial states for benchmarking calculations or simple experimental setups for which histogram-based flux descriptions and simple geometry descriptions in terms of a target mix are adequate. Experiment-specific event generation applications, such as *gT2Kevgen* and *gNuMIevgen*, employ the detailed JPARC and NuMI beam-line simulations and the ROOT-based detector geometry descriptions of the corresponding experiments and are used by a large fraction of the GENIE user base.

On a MacBookPro running MAC OS X 10.5.6, with a 2.16 GHz Intel Core 2 Duo and 1 GB DDR2 SDRAM at 667 MHz, and with all event generation threads enabled, GENIE simulates around 70 events/sec for a $\nu_\mu + Fe^{56}$ initial state at $E_\nu = 1GeV$. The speed is 5 events/sec with the detailed nd280 ROOT-based detector geometry description (40 nuclear targets) and the detailed JNUBEAM-based JPARC neutrino beam simulation (4 neutrino species).

6.5. Utilities

Event generation for a realistic experimental setup typically requires of the order of $\sim 10^8 - 10^9$ differential cross section evaluations just in order to select an interaction to be generated. It is therefore very practical to perform the numerical differential cross section integrations at a different stage and save the data for building cross section splines. The event generation components can recycle these splines for performing fast numerical interpolations, greatly improving the event generation efficiency. GENIE provides a utility, *gmkspl*, to generate all required cross section splines for the intended set of neutrino flavors and nuclear targets over the required energy range and write the data in the XML format expected by the event generation components.

For many use-cases it is convenient to analyze the output *GHEP* ROOT event tree and either write-out simpler, flat n-ntuples containing summary information or convert it to a format expected by an experiment-specific application, such as a detector-level simulation that doesn't use the GENIE I/O. GENIE provides an event tree converter, *gntpc* , which writes out a host of alternative plain text, XML or bare-ROOT formats currently is use by GENIE-based applications in client experiments.

Several tools exist for the purposes of validating and tuning the physics models in GENIE. The ultimate source of data for many of these comparisons is the DURHAM Neutrino Scattering Data Resource [114], an online resource with large compilations of neutrino data from many different experiments. Access to this data in the GENIE package is done through the NuValidator [115] program, a GENIE add-on that can be optionally installed. Distributed with the NuValidator in XML format are data from the DURHAM database, electron scattering data from the Jefferson Lab database [112], and publicly available lepton scattering structure function data [139]. Electron scattering data and lepton structure function data are important in evaluating any scheme for handling the perturbative / non-perturbative transition region described in Sec. 3.2.1 [111]. This data is then available for a variety of applications. The NuValidator includes a simple GUI interface allowing one to select data to display together with the GENIE prediction, with full ability to set model parameters to new values. The data can also be accessed for physics model parameter fits possibly including new experimental data as well as external, historical data.

In addition to event generation and validation the flexibility of the GENIE framework also simplifies many downstream analysis tasks. The evaluation of generator-related systematic errors can be greatly simplified through event reweighting [140]. As mentioned in Sec. 5.2, GENIE's ability to allow algorithm configuration using local pools, redundancy of information between

GHEP event records and *Interaction* objects, and ability to identify algorithms/configurations in generated output all support the development of experiment-specific event reweighting programs.

Details on the application described above and on a host of other utilities may be found in the GENIE Physics and User Manual [110].

7. Collaboration Organization

A significant organizational challenge for the GENIE project is defining the working relationship with experimental collaborations. Having many experiments use the same neutrino event generator is a new development in this field and the exact nature of these working relationships will evolve over time. However some realities and goals are clear.

Good two-way communication is essential both for immediate issues like bug reporting and support and for longer term issues like planning of upcoming physics releases. Experiments often set target dates to begin production of Monte Carlo samples based on publication plans and expectations for data-taking. Meeting these deadlines is often a high priority for these collaborations. Since the time-scale for production of large Monte Carlo samples is roughly similar to the timescale for production of new GENIE physics releases, it will be desirable to synchronize, or at least be cognizant of, upcoming experimental deadlines to as large a degree possible. Discussions with collaborations will also focus on priorities for physics model improvements.

One conduit for these collaborations will be through experimental liaisons who serve as the main contact within an experiment on GENIE issues. This person can then report on the experiment's experiences and deadlines, and can help present the experiment's priorities for model improvements to the GENIE collaboration for evaluation. When effort within the GENIE collaboration to incorporate new model work is not available, these liaisons can assist in providing and organizing effort from their collaborations.

Physics model development is partitioned into subtasks including the cross section model, the hadronization model, and the intranuclear rescattering model. These components are all relatively self-contained and have validation procedures independent of the others. Overall tuning and validation of a production release is the responsibility of a separate working group of the collaboration. This task is undertaken on a roughly yearly timescale. This exercise finalizes the model set and determines values of all parameters based on fits to external data. This process also determines parameter errors which can then be used by experiments in the evaluation of generator-related systematic errors.

A default tune - a self consistent set of physics models and parameter values - will be specified by the collaboration for every major release, along with information about the uncertainties on model parameters and possible correlations. It is possible for experiments, using the validation and parameter fitting tools that are provided as part of the GENIE package, to have their own 'tuning' of the generator. This would be desirable from an experiment's perspective as they have access to new data that, due to the aforementioned uncertainties, will probably not be in complete agreement with the default tuning. It is hoped that the development of new models by specific collaborations as part of ongoing analyses will be speedily adopted into GENIE for the benefit of other users, though the decision of when to make their new models public will be left to the discretion of the user collaborations. It will, however, be important that model improvements be merged back into the default tuning on a regular basis so as to prevent the fragmentation of the base set of physics models into many different experiment-tuned versions which then evolve independently over many years.

The collaboration organizes its effort internally through occasionall meetings of physics model and tuning/validation working groups, phone meetings, blogs, and web-based document databases. The GENIE web site [1] is the central repository for all information related to the package. An extensive Physics and User Manual [110], a web-based source code Reference Manual [141] as

well as a support mailing list are available to GENIE users. Hands-on tutorials on GENIE have been given on several occasions at meetings in the U.S., Europe, and Japan, and material from these workshops with introductory and advanced tutorial examples are available on the GENIE web site.

8. Code Availability

8.1. Version control and distribution

GENIE is available from its CVS repository hosted at STFC's Rutherford Appleton Laboratory from where one can access the development version and a series of 'frozen' production-quality releases. The repository is physically located on AFS space and, in read-only mode, can be accessed anonymously. Read-write mode access to the code repository is provided to the GENIE collaborators via SSH and public key authentication. A transition to a SubVersion repository is planned for the near future. Up-to-date details on how to access the source code are given at the GENIE web-site [1] and at the Physics and User Manual [110].

8.2. Supported platforms and external dependencies

GENIE is known to build on many platforms, including all popular LINUX distributions and MAC OS X, and has no OS proprietary dependency. As of GENIE v2.5.1, external dependencies for a minimal installation that can be used for physics MC production include the *ROOT Class Libraries* [12], the *LHAPDF parton density function library* [142], the *PYTHIA-6* LUND MC [143] and two fairly common utilities: the *libxml2* XML parser and the *log4cpp* error logger.

8.3. Versioning scheme and release lifetime

GENIE versions are numbered as '$i.j.k$', where i, j and k are the major, minor and revision indices respectively. The corresponding CVS tag is '$R-i_j_k$'. When a number of significant functionality improvements or additions have been made, the major index is incremented. The minor index is incremented in case of significant fixes or minor feature additions. The revision number is incremented for minor bug fixes and updates. Versions with even minor number correspond to validated, physics production releases. Versions with odd minor number correspond to the development version of 'candidate' releases tagged during the validation stage preceding a physics production release. Tagged versions always have an even revision number. Odd revision numbers correspond to the CVS head.

Because of the effort invested by the experimental communities in generating large samples and understanding the impact of the simulation changes into their physics results, physics production releases are nominally supported for a long term of approximately 18 - 24 months.

8.4. License

GENIE is now distributed under the GPLv3 license agreement [144].

9. Conclusions

GENIE provides a modern and versatile platform for a universal, 'canonical' Neutrino Interaction Physics Monte Carlo whose validity will extend to all nuclear targets and neutrino flavors over a wide range of energies from MeV to PeV scales. Currently, it includes state-of-the-art neutrino interaction physics modeling in the few-GeV energy range which is relevant for the current and near future long-baseline precision neutrino experiments using accelerator-made beams.

The software was designed using object-oriented methodologies and developed entirely in C++ over a period of more than three years, from 2004 to 2007. The design of the package decouples the mechanics of event generation from the physics models, providing modularity, extensibility, and flexibility. The package supports the full life-cycle of generator-related activities, from event generation using detailed detector geometries and flux information, to final analysis tasks such as reweighting and systematic error evaluation. The data, programs, and procedures used to validate and tune the package are all distributed with the package itself, allowing users the ability to easily extend the package and evaluate new models.

The project is supported by a group of physicists from all major experiments operating in this energy range, establishing GENIE as a major HEP event generator collaboration. GENIE has already been adopted by many neutrino experiments, including those using the JPARC and NuMI neutrino beamlines, and will be an important physics tool for the worldwide accelerator neutrino program.

Acknowledgements

This work was supported by the UK Science and Technology Facilities Council / Rutherford Appleton Laboratory, the US Department of Energy, the US National Science Foundation, and the Tufts University Summer Scholars Program.

The authors would like to express our gratitude to G.Irwin (Stanford), B.Viren (Brookhaven Lab), S.Kasahara (Minnesota) and N.West (Oxford) for contributing to the early stages of the evolution of GENIE through example, by developing the MINOS offline framework, and through their inputs and criticisms during the GENIE design reviews.

We would also like to thank our MINOS collaborators, in particular R.Gran (UMD), K.Hofmann (Tufts), M.Kim (Pittsburgh), M.Kordosky (UCL), W.A.Mann (Tufts), J.Morfin (Fermilab) and S.Wojcicki (Stanford), for their contributions in the development, tuning and validation of the default set of physics models in GENIE.

We would also like to thank E. Paschos (Dortmund), S. Mashnik (Los Alamos), A.Bodek (Rochester), O.Lalakulich (Dortmund, Giessen) and T.Leitner (Giessen) for providing important models and results. We also express our gratitude and recognition to C.Reed (NIKHEF), K.Scholberg, C.Little, R.Wendell (Duke), A.Habig, R.Schmidt (UMN Duluth) and D.Markoff (NCCU) for their ongoing effort to extend the GENIE validity range down to the MeV energy scale.

Also we would like to thank C.Backhouse (Oxford), S.Boyd (Warwick), J.J. Gomez Cadenas (IFIC Valencia), Y.Hayato (ICRR), J.Holeczek (Silesia), Z.Krahn (Minnesota), H.Lee (Rochester), J.Lagoda (L'Aquila), S.Manly (Rochester), B.Morgan (Warwick), D.Orme (Imperial), G.Perdue (Fermilab), T.Raufer (RAL), D.Schmitz (Fermilab), E.Schulte (Rutgers), J.Sobczyk (Wroclaw), A.Sousa (Oxford), J.Spitz (Yale), P.Stamoulis (Athens), R.Tacik (Regina), H.Tanaka (UBC), I.Taylor (Imperial), R.Terri (QMUL), V.Tvaskis (Victoria), Y.Uchida (Imperial), S.Wood (JLAB), S.Zeller (LANL), L.Zhu (Hampton) and many others who have used early versions of GENIE for their help in improving the build system, fixing bugs, and contributing comments and tools.

References

[1] C. Andreopoulos, et al., http://www.genie-mc.org.

[2] S. Agostinelli, et al., GEANT4: A simulation toolkit, Nucl. Instrum. Meth. A506 (2003) 250–303.

[3] M. Bahr, et al., Herwig++ Physics and Manual, Eur. Phys. J. C58 (2008) 639–707.

[4] T. Sjostrand, S. Mrenna, P. Skands, A Brief Introduction to PYTHIA 8.1, Comput. Phys. Commun. 178 (2008) 852–867.

[5] H. R. Gallagher, Measurement of the atmospheric neutrino flavor ratio in Soudan 2, Nucl. Phys. Proc. Suppl. 66 (1998) 290–293.

[6] D. G. Michael, et al., Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam, Phys. Rev. Lett. 97 (2006) 191801.

[7] D. L. Wark, The T2K experiment, In *Thomas, J.A. (ed.) et al.: Neutrino oscillations* 197-215.

[8] D. S. Ayres, et al., Nova proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the fermilab numi beamline.

[9] E. C. Schulte, The MINERvA experiment, Acta Phys. Polon. B37 (2006) 2379–2386.

[10] M. Soderberg, MicroBooNE: A New Liquid Argon Time Projection Chamber Experiment.

[11] M. Soderberg, ArgoNeuT: A Liquid Argon Time Projection Chamber Test in the NuMI Beamline.

[12] R. Brun, F. Rademakers, ROOT: An object oriented data analysis framework, Nucl. Instrum. Meth. A389 (1997) 81–86.

[13] H. Gallagher, The neugen neutrino event generator, Nucl. Phys. Proc. Suppl. 112 (2002) 188–194.

[14] D. A. Harris, et al., Neutrino scattering uncertainties and their role in long baseline oscillation experiments, hep-ex/0410005.

[15] F. Cavanna, O. Palamara, Geneve: A monte carlo generator for neutrino interactions in the intermediate-energy range, Nucl. Phys. Proc. Suppl. 112 (2002) 183–187.

[16] Y. Hayato, Neut, Nucl. Phys. Proc. Suppl. 112 (2002) 171–176.

[17] D. Casper, The nuance neutrino physics simulation, and the future, Nucl. Phys. Proc. Suppl. 112 (2002) 161–170.

[18] A. Rubbia, Nux-neutrino generator, Talk at the 1st Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region (NuINT01) http://neutrino.kek.jp/nuint01/slide/Rubbia.1.pdf (2001) 29.

[19] C. Juszczak, J. A. Nowak, J. T. Sobczyk, Simulations from a new neutrino event generator, Nucl. Phys. Proc. Suppl. 159 (2006) 211–216.

[20] T. Leitner, L. Alvarez-Ruso, U. Mosel, Charged current neutrino nucleus interactions at intermediate energies, Phys. Rev. C73 (2006) 065502.

[21] A. Fasso, et al., The physics models of FLUKA: Status and recent development.

[22] T. Ishida, Charged-current inclusive distributions from K2K near detectors,Prepared for 1st Workshop on Neutrino - Nucleus Interactions in the Few GeV Region (NuInt01), Tsukuba, Japan, 13-16 Dec 2001.

[23] A. A. Aguilar-Arevalo, et al., Measurement of muon neutrino quasi-elastic scattering on carbon, Phys. Rev. Lett. 100 (2008) 032301.

[24] F. Sanchez, Search for neutrino-induced charged current coherent pion production with carbon in a 1.3-gev wide band muon neutrino beam, Nucl. Phys. Proc. Suppl. 155 (2006) 239–241.

[25] A. A. Aguilar-Arevalo, et al., Unexplained Excess of Electron-Like Events From a 1-GeV Neutrino Beam, Phys. Rev. Lett. 102 (2009) 101802.

[26] E. A. Paschos, A. Kartavtsev, G. J. Gounaris, Coherent pion production by neutrinos on nuclei, Phys. Rev. D74 (2006) 054007.

[27] S. K. Singh, M. Sajjad Athar, S. Ahmad, Nuclear effects in neutrino induced coherent pion production at k2k and miniboone, Phys. Rev. Lett. 96 (2006) 241801.

[28] L. Alvarez-Ruso, L. S. Geng, S. Hirenzaki, M. J. Vicente Vacas, Charged current neutrino induced coherent pion production, Phys. Rev. C75 (2007) 055501.

[29] A. Bodek, Muon internal bremsstrahlung: a conventional explanation for the excess electron-neutrino events in MiniBooNE, hep-ex/0709.4004.

[30] J. A. Harvey, C. T. Hill, R. J. Hill, Anomaly mediated neutrino-photon interactions at finite baryon density, Phys. Rev. Lett. 99 (2007) 261601.

[31] O. Buss, T. Leitner, U. Mosel, L. Alvarez-Ruso, The influence of the nuclear medium on inclusive electron and neutrino scattering off nuclei, Phys. Rev. C76 (2007) 035502.

[32] O. Benhar, N. Farina, H. Nakamura, M. Sakuda, R. Seki, Electron and neutrino nucleus scattering in the impulse approximation regime, Phys. Rev. D72 (2005) 053005.

[33] J. E. Amaro, M. B. Barbaro, J. A. Caballero, T. W. Donnelly, J. M. Udias, Final-state interactions and superscaling in the semi- relativistic approach to quasielastic electron and neutrino scattering, Phys. Rev. C75 (2007) 034613.

[34] A. Bodek, J. L. Ritchie, Further studies of fermi motion effects in lepton scattering from nuclear targets, Phys. Rev. D24 (1981) 1400.

[35] E. J. Moniz, et al., Nuclear fermi momenta from quasielastic electron scattering, Phys. Rev. Lett. 26 (1971) 445–448.

[36] H. De Vries, C. W. De Jager, C. De Vries, Nuclear charge and magnetization density distribution parameters from elastic electron scattering, Atom. Data Nucl. Data Tabl. 36 (1987) 495–536.

[37] A. Bodek, U. K. Yang, Higher twist, ξ_w scaling, and effective LO PDFs for lepton scattering in the few GeV region, J. Phys. G29 (2003) 1899–1906.

[38] H. Ejiri, Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei, Phys. Rev. C48 (1993) 1442–1444.

[39] K. Kobayashi, et al., Detection of nuclear de-excitation gamma-rays in water Cherenkov detector, Nucl. Phys. Proc. Suppl. 139 (2005) 72–76.

[40] C. H. Llewellyn Smith, Neutrino reactions at accelerator energies, Phys. Rept. 3 (1972) 261.

[41] R. G. Sachs, High-Energy Behavior of Nucleon Electromagnetic Form Factors, Phys. Rev. 126 (1962) 2256–2260.

[42] H. Budd, A. Bodek, J. Arrington, Modeling quasi-elastic form factors for electron and neutrino scattering, hep-ex/0308005.

[43] R. Bradford, A. Bodek, H. S. Budd, J. Arrington, A new parameterization of the nucleon elastic form factors, Nucl. Phys. Proc. Suppl. 159 (2006) 127–132.

[44] L. A. Ahrens, et al., Measurement of Neutrino - Proton and anti-neutrino - Proton Elastic Scattering, Phys. Rev. D35 (1987) 785.

[45] D. Rein, L. M. Sehgal, Neutrino excitation of baryon resonances and single pion production, Ann. Phys. 133 (1981) 79.

[46] R. P. Feynman, M. Kislinger, F. Ravndal, Current matrix elements from a relativistic quark model, Phys. Rev. D3 (1971) 2706–2732.

[47] K. S. Kuzmin, V. V. Lyubushkin, V. A. Naumov, Axial masses in quasielastic neutrino scattering and single-pion neutrinoproduction on nucleons and nuclei, Acta Phys. Polon. B37 (2006) 2337–2348.

[48] D. Rein, L. M. Sehgal, Coherent π^0 production in neutrino reactions, Nucl. Phys. B223 (1983) 29.

[49] W. M. Yao, et al., Review of particle physics, J. Phys. G33 (2006) 1–1232.

[50] D. Rein, L. M. Sehgal, PCAC and the deficit of forward muons in π^+ production by neutrinos, Phys. Lett. B657 (2007) 207–209.

[51] L. W. Whitlow, S. Rock, A. Bodek, E. M. Riordan, S. Dasu, A Precise extraction of R $= \sigma_L/\sigma_T$ from a global analysis of the SLAC deep inelastic e p and e d scattering cross-sections, Phys. Lett. B250 (1990) 193–198.

[52] M. Gluck, E. Reya, A. Vogt, Dynamical parton distributions revisited, Eur. Phys. J. C5 (1998) 461–470.

[53] S. G. Kovalenko, Quasielastic neutrino production of charmed baryons from the point of view of local duality, Sov. J. Nucl. Phys. 52 (1990) 934–936.

[54] M. Bischofberger, Quasi elastic charm production in neutrino nucleon scattering, UMI-C821490.

[55] M. A. G. Aivazis, F. I. Olness, W.-K. Tung, Leptoproduction of heavy quarks. 1. General formalism and kinematics of charged current and neutral current production processes, Phys. Rev. D50 (1994) 3085–3101.

[56] G. De Lellis, F. Di Capua, P. Migliozzi, Prediction of charm-production fractions in neutrino interactions. ((U)) ((V)), Phys. Lett. B550 (2002) 16–23.

[57] C. Peterson, D. Schlatter, I. Schmitt, P. M. Zerwas, Scaling violations in inclusive e+ e-annihilation spectra, Phys. Rev. D27 (1983) 105.

[58] P. D. B. Collins, T. P. Spiller, The fragmentation of heavy quarks, J. Phys. G11 (1985) 1289.

[59] D. Y. Bardin, V. A. Dokuchaeva, Muon energy spectrum in inverse muon decay, Nucl. Phys. B287 (1987) 839.

[60] W. J. Marciano, Z. Parsa, Neutrino-electron scattering theory, J. Phys. G29 (2003) 2629–2645.

[61] K. S. Kuzmin, V. V. Lyubushkin, V. A. Naumov, How to sum contributions into the total charged-current neutrino nucleon cross section, hep-ph/0511308.

[62] D. MacFarlane, et al., Nucleon Structure Functions from High-Energy Neutrino Interactions with Iron and QCD Results, Z. Phys. C26 (1984) 1–12.

[63] J. P. Berge, et al., Total neutrino and anti-neutrino charged current cross section measurements in 100 GeV, 160 GeV and 200 GeV narrow band beams, Z. Phys. C35 (1987) 443.

[64] S. Ciampolillo, et al., Total cross section for neutrino charged current interactions at 3 GeV and 9 GeV, Phys. Lett. B84 (1979) 281.

[65] D. C. Colley, et al., Cross sections for charged current neutrino and anti-neutrino interactions in the energy range 10 GeV to 50 GeV, Zeit. Phys. C2 (1979) 187.

[66] P. Bosetti, et al., Total cross sections for muon-neutrino and anti-muon- neutrino charged current interactions between 20 GeV and 200 GeV, Phys. Lett. B110 (1982) 167.

[67] A. I. Mukhin, et al., Energy dependence of total cross section for neutrino and anti-neutrino interactions at energies below 35 GeV, Sov. J. Nucl. Phys. 30 (1979) 528.

[68] D. S. Baranov, et al., Measurements of the $\nu_\mu N$ total cross section at 2 GeV - 30 GeV in SKAT neutrino experiment, Phys. Lett. B81 (1979) 255.

[69] S. J. Barish, et al., Study of Neutrino Interactions in Hydrogen and Deuterium: Inelastic Charged Current Reactions, Phys. Rev. D19 (1979) 2521.

[70] N. J. Baker, et al., Total cross sections for $\nu_\mu n$ and $\bar{\nu_\mu} p$ charged current interactions in the 7-ft bubble chamber, Phys. Rev. D25 (1982) 617–623.

[71] T. Eichten, et al., Measurement of the neutrino - nucleon anti-neutrino - nucleon total cross sections, Phys. Lett. B46 (1973) 274–280.

[72] W. Lerche, et al., Experimental Study of the Reaction $\nu_\mu p \to \mu^- p \pi^+$. GARGAMELLE Neutrino Propane Experiment, Phys. Lett. B78 (1978) 510–514.

[73] V. V. Ammosov, et al., Study of the reaction $\nu_\mu p \to \mu^- \Delta^{++}$ at energies 3 GeV to 30 GeV, Sov. J. Nucl. Phys. 50 (1989) 67–69.

[74] H. J. Grabosch, et al., Cross section measurements of single pion production in charged current neutrino and anti-neutrino interactions, Z. Phys. C41 (1989) 527.

[75] J. Bell, et al., Cross section measurements for the reactions $\nu_\mu p \to \mu^- \pi^+ p$ and $\nu_\mu p \to \mu^- K^+ p$ at high energies, Phys. Rev. Lett. 41 (1978) 1008.

[76] T. Kitagaki, et al., Charged current exclusive pion production in neutrino deuterium interactions, Phys. Rev. D34 (1986) 2554–2565.

[77] P. Allen, et al., Single π^+ production in charged current neutrino - hydrogen interactions, Nucl. Phys. B176 (1980) 269.

[78] P. Allen, et al., A study of single meson production in neutrino and anti-neutrino charged current interactions on protons, Nucl. Phys. B264 (1986) 221.

[79] D. Allasia, et al., Investigation of exclusive channels in neutrino / anti-neutrino deuteron charged current interactions, Nucl. Phys. B343 (1990) 285–309.

[80] J. Campbell, et al., Study of the reaction $\nu_\mu p \to \mu^- \pi^+ p$, Phys. Rev. Lett. 30 (1973) 335–339.

[81] G. M. Radecky, et al., Study of single pion production by weak charged currents in low energy neutrino - deuterium interactions, Phys. Rev. D25 (1982) 1161–1173.

[82] D. Day, et al., Study of neutrino - deuterium charged current two pion production in the threshold region, Phys. Rev. D28 (1983) 2714–2720.

[83] T. Yang, C. Andreopoulos, H. Gallagher, P. Kehayias, A hadronization model for the MINOS experiment, AIP Conf. Proc. 967 (2007) 269–275.

[84] D. Zieminska, et al., Charged particle multiplicity distributions in neutrino n and neutrino p charged current interactions, Phys. Rev. D27 (1983) 47–57.

[85] Z. Koba, H. B. Nielsen, P. Olesen, Scaling of multiplicity distributions in high-energy hadron collisions, Nucl. Phys. B40 (1972) 317–334.

[86] G. T. Jones, et al., Neutral strange particle production in neutrino and anti- neutrino charged current interactions on protons, Z. Phys. C57 (1993) 197–210.

[87] N. J. Baker, et al., Strange particle production in neutrino - neon charged current interactions, Phys. Rev. D34 (1986) 1251–1264.

[88] D. DeProspo, et al., Neutral strange particle production in neutrino and anti-neutrino charged current interactions on neon, Phys. Rev. D50 (1994) 6691–6703.

[89] P. Bosetti, et al., Strange particle production in neutrino and anti-neutrino neon interactions, Nucl. Phys. B209 (1982) 29.

[90] M. Derrick, et al., Properties of the Hadronic System Resulting from anti- Muon-neutrino p Interactions, Phys. Rev. D17 (1978) 1.

[91] A. M. Cooper-Sarkar, Hadron final state results from BEBCInvited talk presented at Neutrino '82, Balatonfured, Hungary, Jun 14-19, 1982.

[92] A. B. Clegg, A. Donnachie, A description of jet structure by p_T limited phase space, Zeit. Phys. C13 (1982) 71.

[93] J. A. Nowak, J. T. Sobczyk, Hadron production in Wroclaw neutrino event generator, Acta Phys. Polon. B37 (2006) 2371–2378.

[94] P. Allen, et al., Multiplicity distributions in neutrino - hydrogen interactions, Nucl. Phys. B181 (1981) 385.

[95] D. Allasia, et al., Fragmentation in neutrino and anti-neutrino charged current interactions on proton and neutron, Z. Phys. C24 (1984) 119–131.

[96] A. A. Ivanilov, et al., Inclusive γ and π^0 production in neutrino A and anti-neutrino A interacions up to 30 GeV, Yad. Fiz. 41 (1985) 1520–1534.

[97] H. Grassler, et al., Multiplicities of secondary hadrons produced in neutrino p and anti-neutrino p charged current interactions, Nucl. Phys. B223 (1983) 269.

[98] J. P. Berge, et al., Inclusive Negative Hadron Production from High-Energy anti-neutrino Nucleus Charged Current Interactions, Phys. Rev. D18 (1978) 3905.

[99] V. Ammosov, et al., A study of semi-inclusince gamma production in charged current anti-neutrino - nucleon interactions, Nuovo Cim. A51 (1979) 539.

[100] D. S. Baranov, et al., An estimate for the formation length of hadrons in neutrino interactions, PHE 84-04.

[101] S. Dytman, Final State Interactions in Neutrino-Nucleus Experiments, Acta Phys. Polon. B40 (2009) 2445.

[102] R. Merenyi, et al., Determination of pion intranuclear rescattering rates in ν_μ Ne versus ν_μ D interactions for the atmospheric neutrino flux, Phys. Rev. D45 (1992) 743–751.

[103] R. A. Arndt, W. J. Briscoe, I. I. Strakovsky, R. L. Workman, M. M. Pavan, Dispersion Relation Constrained Partial Wave Analysis of πN Elastic and $\pi N \to \eta N$ Scattering Data: The Baryon Spectrum, Phys. Rev. C69 (2004) 035213.

[104] S. G. Mashnik, A. J. Sierk, K. K. Gudima, M. I. Baznat, CEM03 and LAQGSM03: New modeling tools for nuclear applications, J. Phys. Conf. Ser. 41 (2006) 340–351.

[105] S. G. Mashnik, A. J. Sierk, Recent Developments of the Cascade-Exciton Model of Nuclear Reactions, J. Nucl. Sci. Tech. S2 (2002) 720–725.

[106] D. Ashery, et al., True absorption and scattering of pions on nuclei, Phys. Rev. C23 (1981) 2173–2185.

[107] B. W. Allardyce, et al., Pion reaction cross-sections and nuclear sizes, Nucl. Phys. A209 (1973) 1–51.

[108] A. S. Carroll, et al., Pion-Nucleus Total Cross-Sections in the (3,3) Resonance Region, Phys. Rev. C14 (1976) 635–638.

[109] A. S. Clough, et al., Pion-Nucleus Total Cross-Sections from 88-MeV to 860- MeV, Nucl. Phys. B76 (1974) 15.

[110] C. Andreopoulos, others (The GENIE collaboration), GENIE Physics and User Manual, GENIE-doc-753 (http://projects-docdb.fnal.gov/cgi-bin/DocumentDatabase).

[111] H. Gallagher, Using electron scattering data to tune neutrino Monte Carlos, AIP Conf. Proc. 698 (2004) 153–157.

[112] S. Wood, http://hallcweb.jlab.org/resdata.

[113] D. Bhattacharya, Neutrino and antineutrino inclusive charged-current cross section measurement with the MINOS near detector, FERMILAB-THESIS-2009-11.

[114] M. R. Whalley, A new neutrino cross section database, Nucl. Phys. Proc. Suppl. 139 (2005) 241–246.

[115] C. Andreopoulos, H. Gallagher, Tools for neutrino interaction model validation, Nucl. Phys. Proc. Suppl. 139 (2005) 247–252.

[116] P. Adamson, et al., A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam, Phys. Rev. D77 (2008) 072002.

[117] M. Derrick, et al., Hadron production mechanisms in anti-neutrino - proton charged current interactions, Phys. Rev. D24 (1981) 1071.

[118] D. S. Baranov, et al., Properties of multiplicity distributions in $\nu_\mu A$ and $\bar{\nu_\mu} A$ interactions at E <= 30 GeV, Z. Phys. C21 (1984) 189.

[119] T. Yang, C. Andreopoulos, H. Gallagher, K. Hoffmann, P. Kehayias, A Hadronization Model for Few-GeV Neutrino Interactions, Eur. Phys. J. C63 (2009) 1–10.

[120] N. Schmitz, Production of hadrons in charged current neutrino and anti- neutrinos reactionsRapporteur's talk given at 1981 Int. Symp. on Lepton and Photon Interactions at High Energy, Bonn, West Germany, Aug 24-29, 1981.

[121] P. Astier, et al., Search for $\nu_\mu \to \nu_e$ oscillations in the NOMAD experiment, Phys. Lett. B570 (2003) 19–31.

[122] C. Angelini, et al., New experimental limits on $\nu_\mu \to \nu_e$ oscillations, Phys. Lett. B179 (1986) 307.

[123] S. Dytman, H. Gallagher, M. Kordosky, Hadronic Shower Energy Scale Uncertainty in the MINOS Experiment, arXiv:0806.2119.

[124] K. Scholberg, Neutrinos at the Spallation Neutron Source, AIP Conf. Proc. 967 (2007) 72–75.

[125] B. Frois, C. N. Papanicolas, Electron Scattering and Nuclear Structure, Ann. Rev. Nucl. Part. Sci. 37 (1987) 133–176.

[126] O. Benhar, D. day, I. Sick, Inclusive quasi-elastic electron-nucleus scattering, Rev. Mod. Phys. 80 (2008) 189–224.

[127] A. M. Ankowski, J. T. Sobczyk, Construction of spectral functions for medium nuclei, Phys. Rev. C77 (2008) 044311.

[128] C. Andreopoulos, S. Dytman, H. Gallagher, T. Yang, GENIE v2.0.0 Validation - Comparisons with NEUGEN/daikon+, GENIE-doc-755 (http://projects-docdb.fnal.gov/cgi-bin/DocumentDatabase).

[129] R.Hatcher, G.Irwin, S.Kasahara, B.Viren, N.West, et al., Private communication.

[130] R. E.Gamma, R.Helm, J.Vlissides, Design Patterns, Addison-Wesley, 1985.

[131] ISO 8879:1986(E), Standard Generalized Markup Language.

[132] R. Chytracek, J. McCormick, W. Pokorski, G. Santin, Geometry description markup language for physics simulation and analysis applications, IEEE Trans. Nucl. Sci. 53 (2006) 2892.

[133] A.K.Ichikawa, et al., Private communication.

[134] Y. Fukuda, et al., The Super-Kamiokande detector, Nucl. Instrum. Meth. A501 (2003) 418–462.

[135] K. Anderson, et al., The NuMI Facility Technical Design Report, FERMILAB-DESIGN-1998-01.

[136] D. G. Michael, et al., The Magnetized steel and scintillator calorimeters of the MINOS experiment, Nucl. Instrum. Meth. A596 (2008) 190–228.

[137] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, T. Stanev, A three-dimensional calculation of atmospheric neutrinos, Phys. Rev. D70 (2004) 023006.

[138] G. Battistoni, A. Ferrari, T. Montaruli, P. R. Sala, Progresses in the validation of the FLUKA atmospheric nu flux calculations, Nucl. Phys. Proc. Suppl. 110 (2002) 336–338.

[139] T. Gehrmann, R. G. Roberts, M. R. Whalley, A compilation of structure functions in deep inelastic scattering, J. Phys. G25 (1999) A1–A157.

[140] J. Dobson, A. C., Propagating Neutrino Interaction Uncertainties via Event Reweighting, Acta Phys. Polon. B40 (2009) 2613.

[141] C. Andreopoulos, et al., http://doxygen.genie-mc.org.

[142] M. R. Whalley, D. Bourilkov, R. C. Group, The Les Houches Accord PDFs (LHAPDF) and Lhaglue.

[143] T. Sjostrand, S. Mrenna, P. Skands, PYTHIA 6.4 physics and manual, JHEP 05 (2006) 026.

[144] Free Software Foundation (FSF), GNU General Public License (GPL) version 3, http://gplv3.fsf.org/.